

ICASE

SEMIANNUAL REPORT

October 1, 1995 through March 31, 1996

NASA Contract No. NAS1-19480
May 1996

Institute for Computer Applications in Science and Engineering
NASA Langley Research Center
Hampton, VA 23681-0001

Operated by Universities Space Research Association



National Aeronautics and
Space Administration

Langley Research Center
Hampton, Virginia 23681-0001

DTIC QUALITY INSPECTED

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

19960813 131

CONTENTS

	Page
Introduction	ii
Research in Progress	
Applied and Numerical Mathematics	1
Fluid Mechanics	10
Applied Computer Science	20
Reports and Abstracts	34
ICASE Interim Reports	52
ICASE Colloquia	53
ICASE Staff	56

INTRODUCTION

The Institute for Computer Applications in Science and Engineering (ICASE)* is operated at the Langley Research Center (LaRC) of NASA by the Universities Space Research Association (USRA) under a contract with the Center. USRA is a nonprofit consortium of major U. S. colleges and universities.

The Institute conducts unclassified basic research in applied mathematics, numerical analysis and algorithm development, fluid mechanics, and computer science in order to extend and improve problem-solving capabilities in science and engineering, particularly in the areas of aeronautics and space research.

ICASE has a small permanent staff. Research is conducted primarily by visiting scientists from universities and industry who have resident appointments for limited periods of time as well as by visiting and resident consultants. Members of NASA's research staff may also be residents at ICASE for limited periods.

The major categories of the current ICASE research program are:

- Applied and numerical mathematics, including multidisciplinary design optimization;
- Theoretical and computational research in fluid mechanics in selected areas of interest to LaRC, such as transition, turbulence, combustion, and acoustics;
- Applied computer science: system software, systems engineering, and parallel algorithms.

ICASE reports are considered to be primarily preprints of manuscripts that have been submitted to appropriate research journals or that are to appear in conference proceedings. A list of these reports for the period October 1, 1995 through March 31, 1996 is given in the Reports and Abstracts section which follows a brief description of the research in progress.

*ICASE is operated at NASA Langley Research Center, Hampton, VA, under the National Aeronautics and Space Administration, NASA Contract No. NAS1-19480. Financial support was provided by NASA Contract Nos. NAS1-19480, NAS1-18605, NAS1-18107, NAS1-17070, NAS1-17130, NAS1-15810, NAS1-16394, NAS1-14101, and NAS1-14472.

RESEARCH IN PROGRESS

APPLIED AND NUMERICAL MATHEMATICS

SAUL ABARBANEL

Bounded-error algorithms for long-time integrations of PDE's

The SAT methodology, which seeks to give a priori estimates on the error, has been very successful when applied to the multi-dimensional diffusion equation on complex shapes; (see ICASE Report 96-8). This work was conducted in collaboration with A. Ditkowski.

The idea has now been extended to the case of the (linear) advection-diffusion equation in one and two dimensions. In the 1-D case, solving for the steady case with boundary conditions of $u(0) = 1$ and $u(1) = 0$, we can get very good solutions in the whole domain even though the "boundary layer" near $x = 1$ is not resolved. Standard algorithms with central differencing will cause oscillations which grow with the cell Reynolds number. The SAT method, due to the boundedness of the error, gives smooth results. In the 2D case similar results are obtained. As an added bonus the SAT algorithm converges to steady state at least an order of magnitude faster than the corresponding standard scheme. For truly time dependent problems the error indeed remains bounded, while the usual scheme may diverge for a variety of reasons. These results for the advection-diffusion problem will appear shortly in an ICASE report.

It is planned to generalize the SAT methodology to the case of a system of equations.

EYAL ARIAN

Aerodynamic optimization

The objective of this project is to develop efficient numerical methods to solve large scale optimization problems which are governed by flow equations.

The approach is composed of the following principles: 1) Derivation of the adjoint equations in the continuous space followed by discretization, thus avoiding the computation of grid sensitivities. 2) Derivation of a preconditioner using local mode analysis to approximate the symbol of the Hessian. 3) Performing optimization steps on a hierarchy of length scales by a proper application of multigrid methods.

Numerical implementation for a 2D airfoil with Euler flow are in progress using TLNS3D. Future plans include implementation in 3D and treatment of viscous flows. This work is done in collaboration with V. Vatsa of NASA Langley.

Local mode analysis of the Hessian for optimization governed by PDE

This work is concerned with the analysis of the eigenvalue distribution for Hessians resulting from optimization problems governed by PDE. The goal is to construct efficient preconditioners for such problems.

Local mode analysis of the Hessian has been done for optimal shape problems governed by inviscid flow (in collaboration with S. Ta'asan) and for steady state aeroelastic optimization problem.

Future plans include extension of the analysis to viscous flow.

H. THOMAS BANKS

Coupling effects in structural acoustic systems

Several investigations over the last five years have demonstrated the efficacy of employing model-based control techniques to attenuate noise in structural acoustic systems. A basic tenet throughout the development of this methodology has been to employ actuators on the structure and careful modeling of the structural dynamics, acoustic field and structural acoustic coupling mechanisms to reduce structure-borne noise.

In collaboration with Michael Demetriou (Center for Research in Scientific Computation, North Carolina State University) and Ralph Smith (Iowa State University), the mechanisms for coupling in a 3D structural acoustic system have been numerically quantified. This system models an experimental setup constructed in collaboration with R.J. Silcox (Acoustics Division, LaRC) for online structural acoustic control experiments. As part of the current investigation, controller design based *solely* on structural measurements (e.g., output from accelerometers) has been considered. The goal in such designs is to use the structural measurements and *accurate* modeling of the structural acoustic coupling to reconstruct the full acoustic state which is then used for computing the voltages to the controlling patches. The initial investigations have indicated that in many situations, adequate reduction in sound pressure levels can be obtained in this manner with a limited number of structural sensors and no acoustic sensors. The reduction and/or elimination of acoustic sensors in this manner is highly advantageous in interior structural acoustic applications and crucial in many exterior applications such as reduction of noise generated by transformers or submarines.

Upon completion of these numerical studies, we will begin experiments with the modeled structural acoustic system. These experiments will be performed in the Acoustics Division, LaRC.

JOHN A. BURNS

A PDE sensitivity equation method for optimal aerodynamic design

The use of gradient based optimization algorithms in inverse design is well established as a practical approach to aerodynamic design. A typical procedure uses a simulation scheme to evaluate the objective function and its gradient, then passes this information to an optimization algorithm. Once the simulation scheme (CFD flow solver) has been selected and used to provide approximate function evaluations, there are several possible approaches to the problem of computing gradients. One method is to differentiate the simulation scheme and compute design sensitivities that are then used to obtain gradients. Although this approach produces consistent gradients, in shape design applications it requires that mesh sensitivities be computed.

The PDE based approach uses the partial differential equations that define the flow sensitivities as a guide to develop and select efficient computational methods for approximating sensitivities and the corresponding gradients. This approach has the advantage that mesh sensitivities need not be computed in certain shape design problems. Also, it can be shown that proper combinations of discretization schemes produce algorithms with asymptotically consistent gradients so that under

mesh refinement one can establish convergence of the optimal design algorithm. We have developed a framework to show that properly chosen numerical schemes when combined with trust region optimization are convergent. This result has been applied to a number of aerodynamic design problems and we are currently investigating other applications.

Future plans include the extension of these ideas to adjoint methods and the development of convex approximations that produce suboptimal initial designs. This work was done in collaboration with Jeff Borggaard.

JAMES GEER

Singular basis functions with built-in discontinuities

Series expansions of functions, such as Fourier series, perturbation series, etc., are often useful in developing numerical or semi-numerical, semi-analytical algorithms for the solution of differential equations. However, when only a partial sum of such a series is used, some "undesirable" effects (such as the Gibbs phenomena) may be present, or the partial sum may have "difficulty" approximating certain features of the solution, such as boundary layers, internal layers, or various discontinuities.

For a certain class of functions $f(x)$, the combination of a finite sum of certain singular basis functions with "built-in" singularities, along with a finite Fourier series, leads to a sequence of approximations which converges *exponentially* to f in the *maximum norm*, even though f may have a finite number of discontinuities. In particular, these approximations eliminate the Gibbs phenomena. In order to implement this approach, a knowledge of the locations and magnitudes of the jumps in f and its derivatives is necessary. A simple, accurate, and highly robust method of estimating these quantities from a finite number of the Fourier coefficients of f using a least squares technique has been developed and is being analyzed.

The class of basis functions with "built-in" singularities will be extended to include functions that have more general singularities, such as fractional power (e.g., square root) and logarithmic singularities. They will be applied to several model problems which either have discontinuities in the initial data and/or develop discontinuities (or "near-discontinuities") as time increases.

An analysis of an acoustic/viscous splitting technique for computational aeroacoustics

Since Lighthill's original "acoustic analogy" theory of aerodynamic noise generation, much effort has been devoted to the analysis and computation of sound generated by fluid flows. Only recently has enough computing power become available so that it is possible to consider the direct computation of sound fields from the fundamental governing equations, without the artifice of the analog approach.

The underlying philosophy of the approach being considered is that, at low Mach numbers, the sound radiated from a flow can be calculated readily using the viscous, incompressible flow as the forcing function in a linearized analysis. In short, a flow is first computed treating the fluid as *viscous* and *incompressible*, and then this flow is used to "force" a flow in which the fluid is treated as being *inviscid* and *slightly compressible*. An analysis of how these flows can best be coupled is the subject of this study. In particular, for flow with a small Mach number and a large Reynolds number, several versions of the multiple scale perturbation technique are being explored.

Different scalings of several different classes of problems need to be considered and analyzed. In each case, the coupling of the two flows dictated by the perturbation method will need to be analyzed from both a theoretical point of view and from a practical, computational point of view. In particular, the effects of neglecting viscosity in the slightly compressible flow will need to be quantified.

GENE HOU

Structural reanalysis technique for support condition modifications

Structural components are usually fastened or connected together through various support conditions. The addition and removal of such support conditions will certainly change the topology of the structure and result in significant modification of structural performance. Mathematically, these support conditions can be modeled as pointwise constraints and the theorem of Lagrange multipliers can be used to analyze the structure with support condition modifications. This process, however, is tedious and not suitable to accommodate repetitive reanalyses as a new augmented stiffness matrix must be formed and factored for each modified structure. Thus, the main objective of this research is to derive an efficient reanalysis technique to support a design procedure involving support condition modifications.

In this study the Sherman-Morrison identity is first modified and extended to more general cases of matrix modifications. The portion of the augmented stiffness matrix that is related to the support conditions is then considered as the matrix modification. As a result, the modified Sherman-Morrison identity can be directly applied to find the exact solution of the new structure with a modified support condition without reformulating and refactoring its augmented stiffness matrix. The developed structural reanalysis technique has been verified with examples and successfully implemented in MSC/NASTRAN, a commercial finite element code.

Research is underway to extend the developed reanalysis technique to other applications, such as structural modifications in vibration analysis problems.

HIDEAKI KANEKO

Finite element method with singular basis functions

Recently there has been considerable interest in finite element analyses that incorporate singular element functions. A need for introducing some singular elements as part of basis functions in certain finite element analysis arises out of the following considerations. The solution of certain problems, such as a field problem, exhibits highly singular behavior due to geometric features of the spatial domain. It is thought that an incorporation of singular elements that emulate the solution with the standard polynomial elements may perhaps be desirable. In order to make the computations of the finite element method with singular elements more efficient, Hughes and Akin established an algorithm for constructing interpolation functions that have the same interpolation properties of the Lagrange polynomials.

We pointed out in this research that the aforementioned algorithm is sensitive to the locations of the interpolation points that correspond to the singular basis functions. Specifically, we demonstrated numerically that the rate of convergence of a finite element solution varies according to the

locations of these points. A general theoretical explanation is provided for this variance in the rates of convergence.

A further investigation is required toward the establishment of mathematical theory that guarantees the optimal rate of convergence of a finite element method that uses singular functions as part of its basis.

Superconvergence of degenerate kernel method

The degenerate kernel method is a classical method for finding approximate solutions of the second kind Fredholm integral equations ($x - Kx = f$ in operator form). The basic principle of the method is to approximate a kernel, a bivariate function, as a finite sum of univariate functions. An advantage of the method lies in its simplicity, whereas its disadvantage lies in the high cost of computations. If x_n denotes a degenerate kernel approximation, then by the *iterate* of x_n , we mean $x'_n = f + Kx_n$.

We discovered that the rate of convergence of the iterates of degenerate kernel approximations is determined by the method under which the kernel is decomposed. We proved and demonstrated numerically that, when the decomposition is done as a least squares approximation or as an interpolation approximation using a certain set of interpolation points, then the iterates converge twice as fast as the original degenerate kernel solutions provided that the kernel is sufficiently smooth.

To reduce the high computational cost, we propose to introduce a class of wavelets for a decomposition process. We expect to obtain a system of linear equations whose corresponding matrix is sparse as opposed to the normal full matrix that we encounter in the degenerate kernel method.

Superconvergence of collocation method for Hammerstein equations

The Hammerstein equation arises as a reformulation of a class of boundary value problems with nonlinear boundary conditions. The collocation method is one of the widely used numerical methods to approximate the solution of such equations due to its reasonable computation cost.

We proved that the iterates of the collocation solutions for Hammerstein equation converge faster than the original collocation solution, a phenomenon commonly known as a superconvergence. This result extends the results obtained previously for the Galerkin method. The degree of improvement in the rate of convergence depends upon the smoothness of the kernel involved.

MICHAEL LEWIS

Nonlinear programming methods for engineering applications

The goal of this work is to develop nonlinear programming algorithms and optimization methods that are suitable for use in engineering applications, particularly in support of design work.

Some of this work involves the development of nonlinear programming algorithms; in the past year, with Virginia Torczon of the College of William and Mary, I developed extensions of classical pattern search methods for bound constrained problems. I am also looking at effective algorithms for treating large numbers of equality inequality constraints for non-sparse problems. Another direction of this work is to examine how we can effectively use approximation models to reduce computational expense; with Natalia Alexandrov of the NASA LaRC MDO Branch and Virginia Torczon we have developed a suitable theory and are beginning computational tests.

The directions of algorithmic development next to be pursued are the effective active set trust region algorithms and algorithms for problems with state constraints. We also see enormous potential to extending the work on approximation methods to treat constraints as well.

Optimization of systems governed by PDE

This objective of this work is to investigate how we may take advantage of the analysis of infinite-dimensional optimization to develop more effective computational optimization methods.

One thing I have done is to develop an exposition of the calculation of sensitivities for such problems that presents the subject in an orderly manner, from the point of view of nonlinear programming. I have also been developing numerical techniques for parameter estimation based on a nonlinear programming view of domain decomposition. Finally, inspired by work by Ta'asan and Arian, I have developed a set of techniques to approximate Hessians for problems governed by PDE.

The next theoretical question to address is that of a suitable framework for shape optimization. This appears to require certain infinite-dimensional geometric techniques. Another question to be investigated is the structure of Hessians for problems governed by hyperbolic PDE.

JACQUES LIANDRAT

Wavelet based algorithms for numerical analysis

Various numerical methods based on wavelets have appeared for a few years. They are the starting point for new approaches in numerical analysis that could provide very efficient algorithms specially designed for nonstationary and nonlinear problems.

We have continued to improve and test our methods on 2D problems and have addressed the problem of boundary conditions in one dimension. We are currently investigating the implementation of highly adaptive algorithms in connection with efficient approximation of non linear terms in wavelet bases.

Future Work deals with tests in combustion area in 1D or 2D configuration.

DIMITRI MAVRIPLIS

Unstructured mesh techniques for the Navier-Stokes equations

The overall objective of this work is to develop efficient solution procedures for the Navier-Stokes equations on unstructured meshes involving low memory overhead, with eventual application to parallel machines, for both steady and unsteady flows.

Research has been conducted to improve the efficiency of a previously developed unstructured multigrid solver for high-Reynolds number viscous flows. The main impediment to rapid convergence in these cases is due to the high degree of mesh stretching in the wake and boundary layer regions. A semi-coarsening or more generally directional-coarsening technique has been developed to construct "optimal" coarse grids for solving anisotropic problems of this type. The coarsening proceeds by removing neighboring points which are most strongly connected to the current point as determined by the magnitude of the discretization stencil coefficients. Using a second-order upwind scheme, convergence rates of 0.8 per multigrid cycle for two-dimensional steady viscous flow over

a single airfoil have been demonstrated. Another strategy for reducing the cost of large three-dimensional simulations is through the use of adaptive meshes. Adaptive meshing techniques for three dimensional tetrahedral meshes have been implemented using element subdivision concepts. For unsteady flows with bodies in relative motion, moving meshes are required. A mesh movement algorithm has been developed for three dimensional tetrahedral meshes.

The anisotropic multigrid results are very preliminary, and further research is underway to enhance the robustness of this scheme. Once the two-dimensional algorithm has been fully developed, this approach will be extended to three dimensions for use with steady as well as unsteady flows, with static and adaptive meshes.

CHI-WANG SHU

Discontinuous Galerkin method and shock vortex interaction

Our objective is to study and apply high order finite difference, finite elements and spectral methods for problems containing shocks. This will enable us to capture complicated flow structure over long period of time with a relatively coarse grid.

The investigation of the discontinuous Galerkin finite element method, which is carried out jointly with Harold Atkins at NASA Langley, is in the phase of studying the relative efficiency of evaluating the integrations (both over the element and along the face of the element). The traditional approach is to use a Gaussian type numerical quadrature. We have also been testing the idea of using local polynomial expansions and then exact evaluation. Implementation in 1D and 2D of both approaches is underway. Jointly with Gordon Erlebacher and Yousuff Hussaini, we have finished our investigation of shock longitude vortex interaction problem. Many interesting flow structures are found for strong shocks and/or strong vortex.

Research will be continued for high order methods in finite difference, finite elements and spectral schemes.

DAVID SIDILKOVER

New discretizations and fast multigrid solvers for the inviscid flow equations

The main objective of this work is to develop new discretizations for the Euler equations (both incompressible and compressible). The main requirement to these methods is that they should facilitate the construction of an essentially optimal multigrid solver, i.e. whose efficiency is similar to one of a solver for Poisson or Full-Potential equations.

The discrete equations to be solved at each node are "assembled" from the residuals (fluctuations) of the Euler system on the triangles having this node as a common vertex. The assembly of the discrete equations is performed according to a certain "preconditioning" operator. One of the discrete equations obtained this way for the incompressible flow case is the Poisson equation for pressure. Therefore, despite the use of non-staggered grids, there are no unstable modes in pressure. For the compressible case one of the discrete equation obtained using this approach involves the Full-Potential operator acting on pressure. The constructed discretization *implies* the correct derivation and treatment of the additional boundary condition for the pressure. Collective Gauss-Seidel relaxation can be used as a smoother.

Preliminary numerical experiments demonstrate that the convergence rate of a multigrid solver is $\approx .1$, provided the relaxation is performed in the flow direction. The future plans include testing of the developed techniques for the complex geometry problems.

SHLOMO TA'ASAN

Canonical-variables multigrid for inviscid flows

The efficient numerical solution of the Fluid Dynamics equations is a problem of major importance. While solvers for the Euler equations are quite efficient, the Navier-Stokes solvers perform much poorly. This is due to several reasons and was the motivation of a research toward more efficient algorithms.

The fact that Euler solvers for small Mach number perform much worse than a multigrid Poisson solver on the same grid led to use canonical decomposition of the equations into elliptic-hyperbolic parts. This together with different treatment of the hyperbolic and elliptic parts led to Poisson like efficiency for small Mach number flows. The additional use of potentials to represent the velocity enable the construction of new type of schemes for which Poisson operators appear in several places, allowing the use of simple Gauss-Seidel relaxation. The resulting schemes are h-elliptic for the elliptic part of the system, and upwind for the hyperbolic parts. Moreover, unlike previous results of this type which used staggered grids, the new schemes are cell-vertex schemes. Numerical test both in 2D and 3D in nozzles shows that the method perform with Poisson like efficiency for second order schemes.

Future development include flow around airfoils and wings, with a focus on the treatment of the boundary conditions for the rotational part, which is supposed to introduce lift without introducing cuts in the potentials as was done in the past. The treatment of viscous cases requires non-isotropic coarsening and will be the following stage of this research.

V. VENKATAKRISHNAN

Multi-point design for high-lift configurations

A new effort has been initiated in the area of design optimization, with multi-point design as the ultimate goal. Computational Fluid Dynamics has matured to the point that CFD codes can be used to rapidly compute steady state solutions.

The methodology employed is the continuous adjoint approach. The optimization functional is augmented by the constraints, which are the Euler equations, multiplied by the adjoint variables introduced as Lagrange multipliers. A system of linear partial differential equations results for the adjoint variables. The discretization of the adjoint variables is guided by the knowledge gained by examining the discrete adjoint approach. The design parameters include parameters describing the shape, translation and rotation of the elements, and flow parameters, Mach number and angle of attack. The unstructured grid flow solver, FUN2D, is employed for the analysis. The adjoint solver has a similar flavor. Following a shape change, the grid is restructured using a combination of mesh smoothing and edge swapping.

This work is being done in collaboration with W.K. Anderson of NASA Langley Research Center. Preliminary results for single element shape design are encouraging. Research is continuing on multi-element airfoil design.

Higher order discretizations on unstructured grids

The goal of this project is to develop efficient higher order methods on unstructured grids, with particular emphasis on unsteady flows. The plan is to fully understand the capabilities and limitations of the various approaches in one and multiple dimensions for test problems and to extend the methodology to solve the Reynolds averaged Navier-Stokes equations with field equation turbulence models.

A number of ideas from the finite volume and finite element literature are being pursued. These include reconstruction-evolution methods making use of pointwise or cell-averaged data, and the discontinuous Galerkin method. A number of problems have been solved in one dimension for the linear advection equation, the inviscid Burgers equation and the advection-diffusion equation. The extension of some of the above methods to 2D unsteady problems is currently being undertaken.

FLUID MECHANICS

PONNAMPALAM BALAKUMAR

Nonlinear equilibrium solutions in two-phase flows

Currently, the nonlinear instability analysis of two-phase flows is limited to weakly nonlinear theories. Model equations have been derived and analyzed using asymptotic theories. We investigated the nonlinear instability of two phase flows by seeking equilibrium solutions for the Navier-Stokes equations.

We directly solve for the nonlinear equilibrium solutions using a continuation procedure. The nonlinear unstable region is computed for the flow of two immiscible fluids thorough a circular tube.

In the future, we plan to study the secondary instability of these nonlinear equilibrium flows.

ALVIN BAYLISS

Jet noise/structure/flow field interaction at transonic and low supersonic speeds

We consider the interaction between sound from a jet and an array of flexible aircraft-type panels. We consider jets at both transonic and low supersonic speeds. The vibration of aircraft panels in response to this excitation results in structural fatigue and increased interior noise levels. The objectives are to understand the nature of jet noise, the nature of the panel response and resulting radiation, and to determine techniques to control the panel response. This research was done in collaboration with L. Maestrello (NASA) and C.C. Fenno, Jr. (NRC).

We consider a model in which equations for the evolution of sound in a jet (Euler equations) are fully coupled to equations for the panel response and radiation. These equations are solved using a finite-difference scheme which is fourth order in space and second order in time. Mappings are employed to increase resolution of the jet shear layer and nozzle lip region. At present we have completed an analysis of two-dimensional jets accounting for forward motion effects. We are in the process of completing an analysis of the long time response of high subsonic jets, including nearly periodic vortex shedding.

In the future we will investigate the sound from supersonic two-dimensional jets, including both mixing and shock-induced noise sources. We have also developed an axisymmetric version of our code to allow for the simulation of cylindrical nozzles and jets. When this version is fully validated, it will be extended to account for non-axisymmetric effects.

AYODEJI O. DEMUREN

Numerical simulation of complex turbulent jets

Turbulent jets are encountered in many aerodynamic flows and industrial applications. In many of these cases rapid mixing is desirable, either to reduce noise from jet engines, to promote fuel-air mixing in combustion chambers, or to promote rapid dilution of pollutants exhausting into the atmosphere. Jets with non-circular cross-sections tend to experience more rapid spreading and mixing than circular ones. Although it is believed that the complex vorticity field is responsible

for this, the mechanisms are not fully understood. The goal of current research is to increase the understanding of jet mixing processes through direct simulation of jets with elliptical and rectangular cross-sections.

In collaboration with R. Wilson, a third-order Runge-Kutta and a fourth-order compact scheme is utilized to perform the simulation of incompressible turbulent jets with rectangular cross-sections. In order to obtain a good resolution of the evolving jet, a non-uniform expanding grid is utilized. The generality offered by the curvilinear grid makes extension to complex geometries feasible. To satisfy the divergence-free condition for the velocity field, a Poisson equation is solved for pressure which must also be approximated by the same compact scheme for consistency. The accuracy of these methods has previously been demonstrated with several benchmark problems.

Work is now progressing on the simulation of other complex jets and the comparison of the results to experimental data.

GORDON ERLEBACHER

Shock-vortex interactions

The interaction of shocks with longitudinal vortices (whose axis is perpendicular to the shock) has been the subject of relatively few experiments and numerical simulations. However, the propensity of the vortex to break down is enhanced due to the presence of the shock, and this can sometimes lead to stability problems on fighter aircraft. We seek to understand the nonlinear structure of the downstream flow under a wide range of vortex strength and shock Mach numbers. Detailed understanding of this type of the flow will place us one step closer to developing good models for general shock/turbulence interactions.

With both 6th order compact and 3rd order ENO schemes in space, we have applied our shock-fitted code to study the interaction of weak and strong longitudinal vortices with an infinite planar, axisymmetric shock. We have found that in the range of Mach numbers between 1 and 10, there is a linear curve in $(\epsilon M, M)$ space (where ϵ is the vortex circulation and M the shock Mach number, which defines the parameters at which the vortex breaks down. Breakdown occurs more readily for stronger vortices and stronger shocks. For stronger vortex strengths, the strong upstream shock motion near the axis engenders a triple point structure which moves self-similarly away from the axis.

To investigate detonation-turbulence and detonation-vortex interactions, the code will be modified to include reaction via a simple one-step heat release mechanism with a finite-length reaction zone. This research was done in collaboration with C.W. Shu (Brown University) and M.Y. Hussaini (Florida State University).

Transport coefficients in compressible turbulence

This work derives a transport model for weakly compressible turbulence using Yoshizawa's two scale direct interaction approximation. The utility of this formalism in deriving models for incompressible turbulence and for various coupled field problems is well documented. We believe that Yoshizawa's own treatment of compressible turbulence, while suggestive, oversimplifies the problem by introducing decouplings which suppress the possibility of waves and interaction between turbulence and waves. This work is done in collaboration with Robert Rubinstein and Ye Zhou.

By limiting attention to weak compressibility, a more complete treatment of the response matrix required by this theory is possible. The couplings which arise in the response matrix lead to nontrivial physical consequences including modification of the effective speed of sound, an effect predicted on other grounds by Chandrasekhar, an effective turbulent pressure, and an effective bulk viscosity. These effects do not arise in a fully decoupled theory.

The effects of rotation on weakly compressible turbulence will be investigated by the same methods.

MIKHAIL M. GILINSKY

Advanced methods for acoustic and thrust benefits of nozzle aircraft engines

During M. Gilinsky's tenure at the NASA LaRC as a NRC Senior Research Associate (01/10/93–03/18/96), Dr. M. Gilinsky and Dr. John M. Seiner proposed several new methods for improving the mixing of exhaust jets with ambient air. The 3D nozzle design is incorporated by a corrugated cross section nozzle shape with a sinusoidal lip line nozzle edge. Bluebell nozzle designs have shown noise reduction relative to a convergent-divergent round nozzle with design exhaust Mach number $M_e = 1.5$. The best design provides an acoustic benefit near 4dB with about 1% thrust augmentation. The nozzle thrust calculation was based on a full Navier-Stokes equation solver (NSE), and Euler code with boundary layer correction.

We plan to continue the theoretical and experimental research based on these methods for more efficient mixing of a supersonic jet with ambient air, affecting both jet noise reduction and nozzle thrust augmentation. The Bluebell nozzle concept can be applied for either subsonic or supersonic conditions. The top priority research is an application of a Bluebell nozzle concept for the two-contour nozzle with plug, which is the main nozzle design for the future hypersonic or supersonic aircraft engine. The numerical simulation methods can be applied for this problem solution as in the previous cases. Jet noise calculation will be conducted by a simplified approach based on the compressible Rayleigh model and linearized Euler equations. We also propose to develop the previous approaches for further improving nozzle acoustic and thrust benefits using a new Chisel-shaped nozzle, Chisel and Screwdriver-shaped centerbody (plug). All designs are the subject of invention disclosures, submitted to NASA, which pass through the patent procedure.

SHARATH S. GIRIMAJI

Fully explicit, self-consistent algebraic Reynolds stress turbulence model

A model for the anisotropy of Reynolds stress that is fully explicit and self-consistent was developed earlier. A thorough validation of this model in a variety of turbulent flow situations is absolutely essential before it can be used in design codes for computing practical flows.

This algebraic Reynolds stress model was compared against a full Reynolds stress closure model, previous algebraic models and standard $K-\epsilon$ model for a variety of homogeneous flow situations. All the models were evaluated against DNS and/or LES data. The conclusion was that the present model is an excellent substitute for the more expensive full Reynolds stress closure models. The present model has also been compared against other algebraic stress models for a variety of inhomogeneous flow cases. The present model is clearly superior in physical accuracy as well as numerical robustness and computing time.

Further testing and validation need to be performed for complex engineering flows. The indications are that the present model is as robust and computationally inexpensive as the $K-\epsilon$ model, while being more accurate than the other algebraic models.

Algebraic Reynolds stress model for curved flows

Most flows of practical importance are characterized by curved streamlines. In many important cases, accounting for the effect of streamline curvature on the turbulence is crucial. The objective of this project is to develop an algebraic Reynolds stress model that is fully explicit, self-consistent and Galilean invariant.

Such a model has been developed by invoking the weak-equilibrium assumption in an acceleration-based coordinate system that is Galilean invariant. The resulting model is properly sensitive to curvature for circular flows, and appears to be reasonable for non-circular curved flows also.

Thorough validation and testing for non-circular flows is the next step.

Algebraic modeling of thermal flux

Heat transfer is an important ingredient of many practical flows. Modeling the effects of turbulence on heat transfer is the objective of this project.

This project is in its incipient stage. We consider buoyant flows amenable to the Boussinesq approximation as a first step. Strategies aimed towards developing a self-consistent model are now being explored.

The development of a fully explicit, self-consistent model that is also consistent with the algebraic Reynolds stress model is the future plan for this project.

CHESTER E. GROSCH

Mixing enhancement in a hot, supersonic jet

Experimental observations show that the presence of small tabs on the edge of a hot, compressible jet exiting into a slower moving, colder ambient flow can increase the rate of spreading of the jet. This suggests that the rate of mixing of the jet with the ambient fluid is also increased. The objective of this research is to simulate this flow in order to elucidate the physical mechanism responsible for the increased spreading rate and to perform a parameter study in order to understand the effects of varying the flow conditions.

A set of calculations is in progress, using a compressible Navier-Stokes code, to simulate this flow. The calculated flow without the tabs is in good agreement with experimental measurements. Simulations with modeled tabs were also carried out modeling the tabs by pairs of counter rotating vortices. The results of these calculations show that the presence of the tabs increases the spreading rate of the jet. The basic physical mechanism responsible for the enhanced spreading rate was found. In addition, an extensive parametric study brought out the effects of the vortex parameters on the spreading rate.

Further parametric studies over a wide range of convective Mach numbers will be carried out. Other aspects, including the effect of lobes, will also be investigated.

M. EHTESHAM HAYDER

Non-reflecting boundary conditions

Correct treatment of boundary conditions is required for accurate numerical simulations. We are examining and formulating boundary conditions to minimize numerical reflections. This work is being done in collaboration with Fang Hu (ICASE), M.Y. Hussaini (FSU) and Eli Turkel (Tel-Aviv University).

We are examining various buffer layer techniques and conditions based on the asymptotic solutions at the far field for both inflow and outflow. Flow conditions under current investigation include free shear layers, axisymmetric jets, and shock vortex interactions. These techniques are evaluated by numerical solutions to both the linearized and nonlinear Euler and the Navier-Stokes equations.

Numerical simulation of the mixing noise in turbulent flows

Reliable predictive tools are needed to develop noise suppression technology for many practical flows. Mixing noise is an important component of the overall noise in flows such as jets, and shear layers. This work is being done in collaboration with R. Rubinstein and Y. Zhou.

We employ a stochastic approach to simulate the noise source in the linearized Euler equations. The noise source due to turbulence is modeled so that it has the correct inertial range spatial correlation. It has also the correct temporal correlation.

We plan to examine the potential of our approach by comparing numerical simulations with available experiments.

MAURICE HOLT

Review of Godunov methods

Two one-week visits were made to ICASE in the six month period ending 2/29/96, the first in September 1995, the second at the beginning of February 1996. Most of the time was spent on completion of an ICASE report on "Review of Godunov Methods."

In this paper, second order Godunov methods are reviewed. The early versions by Colella and Woodward (PPM) and van Leer (MUSCL) are described in their original form. The simplification of these by Roe, based on approximate Riemann solver, is then presented. Attention is next given to the improvement in MUSCL due to Hancock and van Leer leading to a fuller paper by Huynh. Finally, brief reference is made to TVD and ENO schemes due to Harten.

The remaining time was devoted to a needed revision of an account of Three Dimensional Characteristics Methods to include important Russian contributions by Chushkin, Rusanov and others.

FANG Q. HU

Absorbing boundary condition for Euler equations

The objective of this work is to develop highly accurate numerical boundary conditions suitable for Computational Aero Acoustics.

A new technique has been developed as numerical absorbing boundary condition for the Euler equations in which Perfectly Matched Layer equations for absorbing linear waves are defined. It is shown that the theoretical reflection coefficient is zero at an artificial boundary for acoustic, vorticity and entropy waves, independent of the wave frequency and angle of incidence. Extensions of PML equations to non-uniform flows and nonlinear Euler equations has been carried out. Numerical examples for mixing layers and nonlinear source flow calculations indicate that this is very promising technique applicable to Computational Aero Acoustics as well as steady state CFD calculations. One particular advantage in nonlinear calculations is that it does not require a knowledge of the mean flow at the non-reflecting numerical boundaries and thus the Euler equations do not need to be linearized. In addition, in collaboration with E. Hayder (ICASE) and M.Y. Hussaini (Florida State University), this technique is applied to shock-vortex interactions and jet noise calculations using Navier-Stokes equations and to duct acoustics (J. Manthey, ODU).

Future work will apply this technique to aeroacoustic problems where noise generation is non-linear and compute the sound radiation directly.

Sound radiation of spatially evolving mixing layers by vortex methods

The objective of this work is to develop vortex simulations for computing sound radiation.

In collaboration with J. Martin (Christopher Newport University), the vortex method is applied to a spatially evolving mixing layer. In our approach, the sound radiation from vortex simulations is modeled based on matched asymptotic expansions of the incompressible simulation and compressible far field for small Mach numbers, in which the vortex simulation provides the inner solution and the far-field radiated sound is the outer solution. The two solutions are coupled through a matching process.

Sound radiation directivity will be calculated in a future project. Applications of 3D vortex methods are also planned.

THOMAS L. JACKSON

Algebraic instabilities in boundary layers

Algebraic instabilities arise when some initial disturbances, owing their presence to a finite level of noise present in any flow, grow sufficiently to trigger nonlinear mechanisms or to provide new basic states for secondary instabilities. These instabilities are distinguished from exponential instabilities, where infinitesimal disturbances always grow exponentially in time. The presence of algebraic instabilities may lead to the so-called "bypass mechanisms". Work is continuing on the evolution of disturbances in the Blasius boundary layer flow. This work offers a means whereby completely arbitrary initial input can be specified and the resulting temporal behavior, including both early time transients and the long time asymptotics, can be determined. This work is conducted in collaboration with R. Joslin (NASA), W. Criminale (University of Washington) and D.G. Lasseigne (Old Dominion University).

The bases for the analysis are: (a) linearization of the governing equations; (b) Fourier decomposition in the plane and streamwise direction of the flow; and (c) direct numerical integration of the resulting partial differential equations. The results provide explicitly both the early time transients and the long time asymptotic behavior of any perturbation. With this knowledge it is then

possible to devise means for flow control and it is possible to either delay or enhance disturbances as the need may be. In addition, the important problem of receptivity can also be analyzed within this framework. All linear results are compared to the equivalent spatial problem using DNS.

Future plans include the investigation of algebraic instabilities in viscous pipe flow, as well as the concept of absolute/convective instabilities.

JOE MANTHEY

Numerical methods for computational aeroacoustics

Numerical schemes for computational aeroacoustics are studied for application to duct acoustics. Special issues are the implementation of the boundary conditions at the duct walls and out flow boundary conditions. Many existing higher order finite-difference schemes are not time stable and hence are unsuitable for long time integration.

Eigenvalue stability analysis has been performed for high-order explicit as well as compact implicit schemes with the physical boundary conditions applied. We find that numerical damping is necessary for stability of explicit schemes.

Future work is to apply the acoustically treated liner condition at duct walls and study its implementation in the time-domain along with its associated numerical stability.

JAMES E. MARTIN

Acoustic analysis of a flap edge model

Sound generated at the side edges of flaps is a very important, in some cases the most intense, source of airframe noise. Recently, Sen of the Boeing Company proposed a new physical mechanism for the flap-edge noise source. Sen's model suggests that the flap edge vortex itself can be excited into periodic oscillations as it is perturbed by a small secondary vortex or turbulent eddy. Together with J.C. Hardin (NASA LaRC), analyses of Sen's simple two-dimensional model were performed and an acoustic code developed to study the potential noise production of the proposed mechanism.

In our study, the side edge of a flap is modeled as a rectangular wedge in the presence of which there exists a potential flow as well as a vortex to represent the flap-edge vortex. The Ffowcs Williams-Hawkings equation is integrated over the surface of the flap to determine the acoustic pressure at far-field locations. Noise generation is found to depend upon a characteristic parameter which can be related to typical aircraft parameters. Calculations were performed to determine the effect of neglecting the retarded time difference in the FWH equations and to determine the influence of integrating over different chord and span lengths. The results of this study indicate that the intensity of sound can be reduced by reducing the chord or increasing the thickness of the flap.

In the future, a fully three-dimensional flap edge vortex will be incorporated to consider the analogous 3D setting.

J. RAY RISTORCELLI

Compressible turbulence modeling

Primary concern is with the creation of a consistent set of models for the effects of compressibility. Focus has been on the reversible transfer due the pressure-dilatation as well as the reduced shear anisotropy which appear to be the most important mechanism in flows with important velocity gradients and nominal transport effects. Less effort was expended on the compressible dissipation due to its apparent lack of importance in high R_t and low M_t flows with important turbulence production mechanisms.

The theoretical aspects of the effects of the 1) pressure-dilatation are complete. An algebraic stress closure has been derived using Girimaji's extension of the ASM theory. The results of the analysis has been tested in a mixing layer and a sizeable decrease in mixing layer growth are found. Additional and larger decreases in the mixing layer growth are expected to come from the reduction in the 2) shear stress anisotropy. Theoretical progress has been made at the second-order closure level to obtain the effects of compressibility on the shear anisotropy. 3) In flows with large mean density gradients the usual eddy viscosity/gradient transfer models for the turbulent transport in the ϵ equations are fundamentally wrong. An analysis carefully accounting for the distinction between Favre and Reynolds averaged quantities in the moment equations has identified and resolved the problem. Computations by J. Morrison for a Mach 8 boundary layer have indicated modest improvements. 4) Compressible data bases are typically started with initial conditions that do not reflect finite Mach number effects. An analysis indicating the proper initial conditions for compressible DNS is finished. Further work in this area is being undertaken with G. Blaisdell.

Future plans in this area include the investigation of the practical and computational implications of the analytical work. These issues are being explored in collaboration with S. Thangam. A substantial portion of the work will be theoretical — at the second-order closure level to account for reduction of the shear anisotropy by compressibility.

Aeroacoustics

A pseudo sound analysis was used to obtain the variance of the dilatational field from $k - \epsilon$ and second-order closures. This allows prediction of the near field sound source intensity from RANS type turbulence models. This work continues in conjunction with RANS calculation done by the Aerodynamic and Acoustics Methods Branch, J. Thomas/ M. Macaraeg. Higher order corrections of the theory are expected.

ROBERT RUBINSTEIN

Dissipation rate transport equations

The experimental finding that rotation suppresses the decay of isotropic turbulence is generally reflected in single point models through modifications of the dissipation rate transport equation. Since rotation does not appear explicitly in the exact transport equation for dissipation rate, the required modifications are often introduced in an ad-hoc fashion. A similar problem arises for the dissipation rate equation in free convection, in which a gravity-dependent term not suggested by the exact transport equation is often included to improve model predictions. The present work

attempts to compute the required corrections theoretically. This work is done in collaboration with Ye Zhou (ICASE).

The approach is based on Kraichnan's formulas for inertial range spectral transport power. By substituting appropriate expressions for the correlation functions which appear in these formulas, rotation and gravity dependence emerge naturally. The appropriate model constants are also expressed as convergent integrals which should suggest at least the order of magnitude of the required corrections.

Current calculations have been limited to the destruction terms. Future work will evaluate the production and diffusion terms by using the methods of Yoshizawa's two scale direct interaction approximation.

Sound generation by isotropic and shear turbulence

Time correlations play a key role in sound generation by turbulence since they determine the distribution of radiated acoustic power over frequency. Previous work on this problem has been based on models for the space-time correlation function which incorrectly assume a separation of variables form for this function. This work is done in collaboration with Ye Zhou.

Models for the space-time correlation function of turbulence developed by Tennekes, Kraichnan, and others are used to compute the frequency spectrum of sound radiated by isotropic turbulence. This investigation completes our previous calculation of the total acoustic power. We also compute the space-time correlation function for shear turbulence and compute the shear dependent contribution to the acoustic power spectrum.

This work is being extended to inhomogeneous flows including jets and to the sound radiated by heated jets.

YE ZHOU

Reduced tensors for stress and dissipation

Fluid dynamics needs very intensive computations, and saving of the computational work at each time step can be very useful.

In this theoretical paper we systematically explore the concept and application of reduced stress tensor and dissipation function, showing that they may bring significant savings in analysis and computation. We develop a methodology to ensure that the stress tensor, regardless of its number of independent components, can be reduced to an equivalent one which has the same number of independent components as a surface force. It applies to both the momentum and the energy equation where the shear viscosity is constant and a Reynolds stress model is given. Direct application of this method leads to a reduction of the dissipation rate of kinetic energy. Following this procedure, a significant saving in analysis and computation may be achieved. Moreover, for turbulent flows, we show how the modeling of Reynolds stress tensor can be reduced to that of the mean turbulent Lamb vector. This insight leads us to suggest that modeling the Lamb vector, instead of Reynolds stress tensor, may be more beneficial. As a first step of this model development, we derive the governing equations for the Lamb vector and its square. We believe these equations can form a basis of new second-order closure schemes and should be compared to the traditional Reynolds stress transport equation.

A numerical examination of the budget and spectra of terms in this equation, including those that need be modeled, is being undertaken and will be reported separately. This work was done in collaboration with Drs. James Wu and J.Z. Wu of The University of Tennessee Space Institute.

APPLIED COMPUTER SCIENCE

ABDELKADER BAGGAG

Iterative methods for solving linear systems arising from CFD problems

The last decade has witnessed an explosion of research activities in the area of parallel numerical algorithms. Highly reliable and efficient iterative solvers have been developed for symmetric positive definite linear systems. For symmetric indefinite or general nonsymmetric systems, the most widely used class of Krylov-based methods often fail mainly due to lack of robustness and lack of effective preconditioners.

Realistic systems such as those arising from solving Navier-Stokes problems, via certain discretization schemes, have the form:

$$\begin{bmatrix} A & B \\ B^T & 0 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} f \\ g \end{bmatrix} \quad (1)$$

where A is a symmetric positive definite $n \times n$ matrix, B is an $m \times n$ full rank matrix. This system is generally very large and sparse, and iterative schemes are usually the only feasible methods for obtaining the solution.

There are many approaches to the iterative solution of such systems. We will simply mention some of them here. The modified Conjugate Gradient algorithm may be the easiest one which does not take into account the specific structure of the coefficient matrix. Bramble et al. reformulate the system into a symmetric positive definite one with a well-conditioned matrix. While the theory is quite elegant, it is difficult to construct the required preconditioner, and the observed performance is not competitive. Wittum right-preconditions the system such that the resulting system is suitable for ILU preconditioning. This approach is based on the multigrid schemes which are widely used in the solution of such systems.

Among the more classical methods, the Uzawa-type methods, solve first for the pressure unknowns, which in turn involves solving positive definite systems of the form $A^{-1}x$ exactly or inexact, at each step. Another class of iterative methods is that of projection in which velocity unknowns are computed first and a least square problem is then solved for the pressure unknowns.

Early numerical experiments show applicability of the above methods to the problem at hand. Studies will continue on modifying these schemes to address those numerical simulations of solid-liquid flows. Future work concerns the development of the above schemes on parallel computers and measurement of their performance. Also, study and implementation of the so-called Newton-Krylov-Schwarz will be pursued in close collaboration with Dr. David Keyes at ICASE.

XIAO-CHUAN CAI

Newton-Krylov-Schwarz algorithms for transonic flows

We continue to study Newton-Krylov-Schwarz algorithms (NKS) for the numerical simulation of transonic flows on parallel computers, such as the IBM SP2 at NASA Langley.

NKS employs an inexact finite-difference Newton method and a Krylov subspace iterative method, with a two-level overlapping Schwarz method as a preconditioner. We solve a 2D transonic full potential equation, discretized with a bilinear finite element method and a density upwind scheme used in the elements where the flow is transonic. For subsonic problems, the theoretically expected performance of the method is essentially achieved. For the transonic case, the numerics are more encouraging than existing theory. Overall computation time is approximately six times greater for the transonic than for the subsonic case, with current upwinding strategies. This can be factored into a three-fold increase in the number of Newton steps in the transonic case, and a two-fold increase in the number of Krylov iterations per Newton step. This work has been done in collaboration with David Keyes.

We plan to further study this class of NKS algorithms for more complicated flows that have to be modeled by a combined full potential/Euler scheme.

GIANFRANCO CIARDO

Kronecker operators for the description and solution of Markov models

The solution of large Markov models is limited by the size of the transition rate matrix R . Even when stored in sparse format, the practical limit for a modern workstation is substantially less than 10^6 states and 10^7 nonzero entries, if we want to avoid the use of virtual memory. Many interesting problems in performance and reliability, however, correspond to much larger models. Our objective is to increase these limits by one order of magnitude.

To achieve this goal, we use Kronecker operators (also known as tensor operators) for the description of R . We decompose the model into n interacting submodels, each described by a (small) transition rate matrix $R^i, i = 1, \dots, n$, and express R as the Kronecker sum of the matrices R^i , plus some Kronecker products of "corrective factor" matrices, also of small size, describing the interactions between submodels. We show how to manage models that have an extensive type of state-dependent behavior, both in the structural (what happens when an event occurs) and stochastic (how long does it take to occur) aspects. Then, we eliminate an important restriction existing in previous work: events corresponding to synchronizations between submodels can now be timed or instantaneous. Finally, we show how the decomposition of a large model into submodels can be carried to the level of individual "local events" and "local states." Thus, there is hope for being able to decide algorithmically how to decompose a model. With an early prototype, we were able to solve models with 2.5 million states and 25 million nonzero entries in a few hours on a workstation with only 90 megabytes of main memory, without paging.

Two further improvements should be considered. First, the efficient solution of large linear systems where the matrix is described as a Kronecker expression is only now receiving attention. We are collaborating with leading researchers on this issue with the goal of achieving both theoretical results and efficient implementations in the near future. Second, we recently explored, with David Nicol, the distributed generation and solution of the state-space. This allows the size of models that can be solved to increase by an order of magnitude, provided we have several workstations, or a multiprocessor system. By combining the two approaches, we envision being able to solve models with 10^8 states and 10^9 nonzero entries using a dozen processors having 128 megabytes of main memory each.

THOMAS W. CROCKETT

Parallel graphics libraries for runtime visualization on massively parallel architectures

Large-scale scientific computations which run on massively parallel computers often generate datasets which are too large to manage conveniently on workstation-class systems. By exploiting the available parallelism to perform tasks such as visualization and graphics, these large datasets can be analyzed in place as they are created, reducing the need to move them elsewhere for post-processing.

Under the auspices of the national High Performance Computing and Communication Program (HPCCP) and NASA's Computational Aerosciences project (CAS), we have been developing parallel rendering algorithms and prototype graphics libraries which give parallel application programs the ability to generate live visual output at runtime. This work has progressed to the point that we now have operational software with sufficient functionality to be useful in a variety of visualization applications. During this reporting period, we have concentrated on preparing the software for release to the HPCCP community, including performance enhancements, bug fixes, and especially the development of documentation. The documentation is being produced in hypertext form suitable for online publication via the World Wide Web.

Future plans include the addition of important features such as transparency, a sphere primitive, and color quantization. There are some remaining algorithmic issues with large numbers of processors which we hope to address, and we plan to utilize our software infrastructure as a base for developing parallel visualization algorithms. We are also exploring possible collaborations with researchers at the Jet Propulsion Laboratory and Los Alamos National Laboratory.

PHILLIP M. DICKENS

Thread-based parallel direct execution simulator

Parallel direct execution simulation is an important tool for performance and scalability analysis of large parallel-message programs executing on top of a presumed *virtual computer*. In this approach the application code is executed directly to determine its run-time behavior, and any references to the presumed virtual computer are trapped and handled by simulator constructs. However, detailed simulation of message-passing codes requires a great deal of computation, and we are therefore interested in pursuing implementation techniques which can decrease this cost. This work is being done in collaboration with Matthew Haines, Piyush Mehrotra and David Nicol.

Our approach is to implement the simulated virtual processors as light-weight threads rather than heavy-weight Unix processes thus reducing both on-processor communication costs and context switching costs. However, this approach brings up the important issue of how to provide the separate address spaces required for each of the simulated virtual processors when threads are generally implemented such that they all share one global address space. Our solution is to implement the virtual processors using the C++ class mechanism which automatically provides the required separate address spaces. We have a working version of our thread-based direct execution simulator implemented on a cluster of Sun workstations. Experimental studies on real-world application codes show that the thread-based simulator offers up to a three fold improvement in performance when compared to the process-based approach.

There are several worthwhile extensions to this project. First, we would like to augment the thread-based version with the ability to predict execution times on both an Intel Paragon and a network of workstations. In the latter case, this would provide the important ability to predict performance of codes executing on a *non-dedicated* workstation cluster. Also, we would like to extend our simulator to handle multiple processes per virtual processor. Finally, we would like to explore the use of our simulator for perturbation free performance profiling.

STEPHEN GUATTERY

Multilevel separator algorithms and applications

Graph separators have important applications in numerical computation for tasks such as partitioning and computing elimination orderings via nested dissection. Since finding optimum separators under useful definitions of "optimum" is NP-complete, separator algorithms are often evaluated on the basis of their performance in a particular application. A recent trend has been the use of multi-level separator algorithms. These algorithms work in a way roughly analogous to multigrid algorithms: a graph is coarsened in several steps to produce a smaller graph, which is partitioned. The graph is then uncoarsened; at each uncoarsening step the lower-level cut is extended to the finer graph, and refinement of the cut may be done. Recent algorithms by Karypis and Kumar, Ashcraft and Liu, and Hendrickson and Rothberg have followed this general model. However, problems can arise: the coarsening process may stall, or the separator produced may be good with respect to the criterion used by the separator algorithm, but poor with respect to the application. Pothén and Kurfert have observed a graph from a linear programming problem for which nested dissection algorithms using multilevel separator algorithms produce extremely poor orderings.

We examined this particular graph and abstracted properties of its structure to generate hypotheses about why the orderings produced are bad. Some of the hypotheses have been tested and verified. This led to a list of properties that should be desirable for multilevel algorithms in general (keeping an eye to parallelization) and to a list of properties desirable for separator algorithms used in ordering algorithms. This research is in collaboration with Alex Pothén.

Further efforts in this area will include implementing multilevel separator algorithms consistent with these properties. This will allow us to test if the resulting algorithms perform better, and to refine the details of the properties. We also hope to formalize the empirical studies discussed above in order to produce test graphs that can be used in assessing specific aspects of algorithm performance.

Eigenvalue bounds for Laplacians and Markov Chain transition matrices

Bounding eigenvalues of Laplacian matrices (representing graphs) and Markov chain transition matrices has numerous applications. For instance, lower bounds on the second-smallest eigenvalue of the Laplacian can be used in the analysis of the behavior of spectral separator algorithms. Bounds on the second-largest eigenvalue of the transition matrices of time-reversible Markov chains can be used to bound the mixing time; this has application in sampling algorithms. These two problems are closely related, and similar approaches can be applied in solving them. Various techniques for computing such bounds have been suggested by Sinclair and Jerrum, Sinclair, Diaconis and Stroock, and Kahale. In general, they involve embedding graphs into the graph representation of the matrix.

We extended a previous clique embedding method for Laplacians (work with Miller and Kahale) to Markov chain matrices. This technique provides better bounds for certain chains. We also showed how to use Poincaré Inequality techniques that work with a specified zero boundary to bound the second smallest Laplacian eigenvalue for graphs subject to odd-even decomposition. Attempts to extend a previous result showing that the clique-embedding lower bound is tight to within a log factor for symmetric (edge-weighted) trees has shown that the proof for the symmetric case does not extend simply to the symmetric case.

Future work will focus on finding a proof for a tightness result for general trees. It will also focus on finding algorithmic applications for these ideas.

JIM E. JONES

Parallel semi-coarsening multigrid methods

Semi-coarsening multigrid algorithms are robust and efficient solvers for problems with anisotropic and highly variable coefficients. In particular, they can often be used as “black-box” solvers for the linear system arising from the discretization of elliptic partial differential equations on logically rectangular grids. They represent an important computational algorithm whose parallel implementation seems well suited for High Performance Fortran (HPF), a FORTRAN based language for parallel computing.

In the past six months, we have written a serial FORTRAN code for the standard semi-coarsening algorithm and, after verifying its performance in the serial environment, have added the HPF code to port it to a parallel computer. The code is now being debugged and analyzed on the IBM SP2 at NASA Ames where we have access to several HPF compilers. We have also begun investigating a new variant of the semi-coarsening algorithm. As part of the standard algorithm, a fine grid problem is reduced to a smaller problem, involving only every other grid line, with half the unknowns as the original problem. The variant we are considering uses two such smaller problems — one involving the even grid lines and the other the odd lines. These two smaller problems can be solved concurrently. We have completed a serial implementation of this variant and are currently analyzing its performance.

We plan to continue our HPF implementations of semi-coarsening multigrid algorithms on the IBM SP2. This project should yield information about the current state of HPF compilers and the language’s suitability for parallel multigrid applications. In addition, it will allow parallel performance comparisons between the standard algorithm and variants.

DAVID E. KEYES

Parallel algorithms of Newton-Krylov-Schwarz type

The crux of Newton-Krylov-Schwarz (NKS) algorithms is a balance of implicit convergence rate and good data locality, which we characterize by the slogan “Think Globally, Act Locally.” Though NKS methods are now widely employed in their matrix-free form, their use is often at one extreme (where parallelism is limited) or the other (where convergence rates are sensitive to granularity) of this balance. This is not surprising, since the best balance is generally problem- and architecture-specific, and not easily codified.

In various ICASE-hosted collaborations we have demonstrated NKS methodology in potential, Euler, Navier-Stokes, and combustion codes. In a problem for which a good coarse operator is known, its use in the Schwarz preconditioner provides a convergence rate nearly independent of mesh density and mesh partitioning granularity. Problems in which good coarse operators are not known tend to have strong intercomponent coupling, rendering their convergence rates intrinsically less sensitive to mesh size and partitioning than "bare" elliptic problems.

We are working on partitionings that preserve the strongest couplings subject to concurrency requirements of practical Schwarz preconditioning. We are also seeking to refine pseudo-transient continuation approaches for steady or multiple time-scale problems.

PDE-oriented problem solving environments

The Portable Extensible Toolkit for Scientific Computing (PETSc) being developed at Argonne National Laboratory is a large and versatile package integrating distributed vectors and matrices in several local sparse storage formats, Krylov subspace methods, preconditioners (including Additive Schwarz), and nonlinear solvers (including Newton-based methods and pseudo-transient continuation). Each software component consists of an abstract interface and one or more implementations using particular data structures. This design leads to layered procedures for the various phases of solving PDEs, with a common interface style for each class of problems. PETSc's environment for rapid algorithm development and prototyping can be turned into a high-performance environment by changing compilation flags, as opposed to source code.

Nilan Karunaratne and Jie Zhang of Old Dominion University and I are working with two of PETSc's developers, Lois Curfman McInnes and Barry F. Smith, to try to ensure that the package will accommodate the needs of computational fluid dynamicists (among others). Our experiences with full potential, Euler, and combustion codes are responsible for many existing and evolving features of the package, which currently runs on more than ten architectures, including ICASE machines and both Langley MPPs.

Though porting an application code with complex data structures to the distributed-memory parallel environment may still carry a cost of multiple person-months, experience shows significant opportunities for code reuse once the discretization and solution processes are disentangled. Our short-term goal is to make this process "routine" for codes based on structured grids. In the long term, we will build on this experience in the unstructured grid setting.

Communication modeling in distributed computing

The "hyperbolic model" for communication costs in multi-layer contended networks is a two-parameter characterization of message service time and a set of rules for combining parameter pairs for individual components or subnetworks into end-to-end parameter pairs. It represents a practical compromise between communication models of greater complexity and fidelity, like the five-parameter LogGP model, and unrealistically ideal models, like the PRAM.

With Ion Stoica and Florin Sultan of Old Dominion University, we recently illustrated the application of the hyperbolic model to four rather different distributed-memory architectures available at NASA Langley: an Ethernet network of workstations (NOW), an FDDI NOW, an IBM SP2, and an Intel Paragon, using common MPI primitives. The parameters of the model are evaluated from simple communication patterns. Then overall communication time estimates, which compare

favorably with experimental measurements, are deduced for the message traffic in a time-parallel multigrid code.

For transformational computing on dedicated systems, for which message traffic is describable in terms of a finite number of regular patterns, the model offers a good compromise between the competing objectives of flexibility, tractability, and reliability of prediction. We plan to demonstrate it on larger CFD codes to explain their network performance.

KWAN-LIU MA

Global and local vector field visualization using enhanced line integral convolution

Visualizing vector field data is challenging because no existing natural representation can visually convey large amounts of three dimensional directional information. In fluid flow experiments, external materials such as dye, hydrogen bubbles, or heat energy are injected into the flow. The advection of these external materials can create stream lines, streak lines, or path lines to highlight the flow patterns. Analogues to these experimental techniques have been adopted by scientific visualization researchers. We have developed methods that integrate local and global visualization techniques to explore three-dimensional vector field data on regular grids. This work has been done in conjunction with H. Shen and C.R. Johnson at the University of Utah.

A technique called Lin Integral Convolution (LIC) uses a one dimensional low pass filter to convolve a white noise texture based on the directional information of the vector field. These methods can successfully illustrate the global behavior of vector fields: however, little or no user probing capability is provided, so specific information about the local behavior of the field is limited. We use the LIC method as the underlying algorithm, but we add two very useful improvements. The first improvement permits local probing by allowing the user to introduce "dye" of various colors into the 2D/3D LIC flow field. The inserted dye propagates through the flow field, highlighting local flow features such as wavefronts, while the standard LIC texture still illustrates the global motion. The second improvement uses the direct volume rendering method to map the LIC textures onto any contour surfaces or onto translucent regions derived from multiple scalar quantities. While the LIC algorithm can be directly applied to three-dimensional vector field, there is no systematic post-processing method available to render the resultant LIC data. We propose a rendering model that can intermix LIC images with other scalar data. This rendering model can also be used to explore the cause-and-effect relationships of multiple scalar fields within the same computation grid.

Future work includes providing more flexible control for modeling the behavior of the dye. For example, currently the life span of the dye is dependent on the convolution length of the LIC kernel. However, it is necessary for the user to control the life span without affecting the convolution length globally. Another extension of this research will be to study the effect of using different filter kernels to control the appearance of the dye. An additional feature that would make the software more useful will provide a natural and effective way for the user to probe the three-dimensional space and inject the dye.

Parallel volume rendering of unstructured data: A distributed approach

Unstructured meshes are used to model scientific and engineering problems with complex geometries; an example is the design of aircraft. The adaptive and irregular nature of the meshes

introduces difficulties to visual interpretations of the solution data. In particular, for large unstructured data, powerful visualization techniques like direct volume rendering are often too expensive to run on a single workstation. The aim of this research is to develop a data-distributed rendering algorithm for parallel volume ray-casting on unstructured meshes. The algorithm makes possible rendering of large data sets that cannot fit into the main memory of a single workstation. Work continued on this research studies how the characteristics of the unstructured meshes, the data partitioning, and transfer functions affect the overall performance of the parallel renderer.

We take a completely distributed approach which distinguishes it from previous efforts done by other researchers. Both the data and the rendering computation are distributed across the available processing nodes. Inter-processor communication is only needed for the image compositing step. Within each processor, ray-casting of local data is performed independent of other processors. Image compositing is overlapped with the ray-casting processes to achieve higher parallel efficiency. The algorithm has been tested on the Intel Paragon, a distributed-memory MIMD parallel computer, using from two to 128 processors. When a volume mesh is partitioned for parallel processing, the area of boundary surface increases and so does the corresponding ray-casting cost, since a ray begins at the boundary surface. Therefore, it is not possible to achieve optimal parallel efficiency. Moreover, both data partitioning and the selection of transfer functions determine the actual computational load on each processor. Test results show that imbalanced load further degrades the rendering efficiency.

The performance studies indicate many opportunities for further optimization of the algorithm and its implementation. In particular, ray-casting overhead and imbalanced load significantly affect the overall performance of the renderer. Dynamic load balancing, particularly by the pool of tasks approach, should be implemented to reduce the typical load imbalance which is above 20%, and at the same time to allow more efficient data exploration spatially. As approaching interactive rendering rates using more processors or more powerful processors, it is then needed to reevaluate both the image-space partitioning and image compositing step as the rendering step become less dominant. Data and image I/O, a frequently ignored problem, must be improved to make the overall rendering process more efficient. One important use of such parallel rendering capability is to support runtime monitoring of numerical simulations running on a parallel computer. Therefore, future work will also focus on supporting runtime visualization, along with the development of a rendering library.

PIYUSH MEHROTRA

Software environment for multimodule parallel applications

Exploiting the multiple levels of parallelism in multimodule applications, such as multidisciplinary design of aircraft, requires software support beyond what is provided by data parallel languages such as High Performance Fortran. We have designed a set of language extensions, called Opus, which provide such support. The central concept in Opus is that of a Shared Abstraction (SDA), a set of data structures along with methods to access the data. SDAs can be internally data parallel whereas control parallelism is represented by multiple SDAs executing asynchronously. This work is being done in collaboration with Matthew Haines, John Van Rosendale, Barbara Chapman and Hans Zima and a graduate student, Dawn Galayda.

In the last few months we have focussed on enhancing the runtime system required for implementing internally data parallel SDAs requiring SPMD support across the processors of a distributed machine. Also, data parallel SDAs require special support for method activation and argument handling. Each SDA has a designated leader, which handles method initiation protocol and then activates its workers. To handle distributed data for arguments, leaders of SDAs initially exchange distribution information and then workers independently generate communication schedules for moving input and output data as required. We are also building a source-to-source translator to translate Opus code into Fortran 90 and HPF code. The major issue here is to provide support for separate address spaces for each SDA.

We plan to complete the implementation of the Opus environment in the coming months and use it for simple and complex test cases that we are generating. We plan to continue evaluating the language design, modifying it as necessary based on our evaluation.

Using thread migration as a mechanism for load balancing

While load balancing in multi-processor systems has received a great deal of attention in the past, there has been little work in utilizing light-weight threads for this goal. Since threads are light-weight, their contexts are small and their migration can be accomplished with minimal overhead. It is our belief that we can implement a migration mechanism for threads that will be both efficient and easy to use. With such an implementation we should be able to apply both existing and new load balancing techniques in an attempt to spread the load of an application evenly over the processors being used by the application. This work is being done in collaboration with Matthew Haines and David Cronk, a VILaP-HPCC graduate student.

We have designed and implemented the Chant runtime system which allows us to exploit light-weight threads in a distributed environment. By design Chant allows threads to exist in a distributed environment and allows the system to keep track of threads in a global manner. In addition, Chant has a remote service request (RSR) server built in which should facilitate our efforts to implement thread migration. Chant's portable and runs on a cluster of workstations, the Intel Paragon, and the IBM SP2. We are currently extending Chant to support efficient thread migration.

One of the major issues in thread migration is the handling of pointers. We are looking at several techniques to solve this problem including the use of specialized pointers in migratable threads. Load balancing has been extensively studied; we will focus on new load balancing techniques which benefit from the underlying thread model. In implementing such new techniques, we will also devise methods for determining both when a thread should be migrated as well as where it should be migrated. We plan to compare the efficiency of our load balancing implementation to that of traditional, non-threaded techniques.

Evaluation of HPF

The stated goal of High Performance Fortran (HPF) was to "address the problems of writing data parallel programs where the distribution of data affects performance." We have been using data parallel codes of interest to NASA to evaluate the effectiveness of the language features of HPF. This work is being done in collaboration with the graduate student Kevin Roe.

We have been examining the performance obtained by current HPF compilers on simple CFD problems. The first problem examined was the acoustic wave propagation. It was converted to

HPF and tested using three different HPF compilers. APR's and PGI's HPF compilers were used on NASA Ames's SP-2 and Digital's HPF compiler was used on the DEC Alpha Farm from NPAC at Syracuse University. The second problem examined was the shock tube or Riemann problem. Two HPF compilers (PGI and APR) were used in testing this code on the IBM SP-2 at NASA Ames. The Total Variation Diminishing (TVD) numerical algorithm present in this code presented a problem because of its variable stencil. The size of the stencil is dependent on the value at each grid point and thus is only available at runtime. The code was written in a couple of different ways, one using conditionals and the other using coefficients, and the performance compared. A scaled-down version of TLNS3D-MB, a multiblock code, was parallelized using HPF. Results for 2, 4 and 8 equally sized blocks showed the current HPF compilers' ability for dealing with embarrassingly parallel problems. The section of the model code that required communication was optimized to minimize the number of communication operations generated by the compilers.

In the next six months we plan to focus on the full TLNSD code converting it to HPF. We will also be experimenting with using a partially mixed MPI and HPF code for TLNS3D wherein the boundary data exchange will be handled by MPI message passing calls and the computations within a block will be handled by HPF.

DAVID NICOL

Utilitarian parallel simulator

One of the principal reasons that parallel discrete event simulation has not made an impact on simulation practice is the difficulty of programming the synchronization correctly. Our objective is to alleviate this problem in the context of simulating parallel computer and communications systems. This research is in collaboration with Philip Heidelberger (IBM).

Our approach recognizes that several successful conservative synchronization algorithms are known to be effective in the context of simulating parallel computer and communication systems. We have developed a software library called U.P.S. (Utilitarian Parallel Simulator) that is used in conjunction with the commercial CSIM simulator package, to provide transparent synchronization and communication. The simulation modeler develops models largely as he would for a serial simulation. By incorporating U.P.S. constructs at the interface between simulation processes, the synchronization and communication activity is carried out automatically. U.P.S. has been implemented on the Paragon, and on the IBM SP-2. Recent work includes creating a C++ version to work on the IBM SP-2, and design of features that support dynamic load balancing.

Future work includes development of a GUI interface, and modifications to support integration with direct-execution based tools.

RITE (Reliability Interface Tool Extension)

One of the principle reasons reliability analysis is treated as an afterthought in system design is that the tools and modeling methodology used is very different from that used to design the functional behavior. Our objective is to address the problem by integrating the reliability analysis directly into the functional design tool.

Our approach is to develop a reliability analysis engine designed specifically for such integration. Our current effort—RITE—is an enhancement of our previous work with the tool REST. RITE

addresses several shortcomings of REST, specifically by using new mathematics for the analysis that correct REST's breakdown on long mission times, and by analyzing multiple mission times concurrently.

Future work includes creating a distributed version of RITE that exploits the inherent parallelism in the analysis.

CAN OZTURAN

Distributed environment for adaptive unstructured meshes

The focus of this project is to develop a general purpose distributed unstructured mesh environment to satisfy the needs of parallel adaptive PDE solvers on two- and three-dimensional complex geometry. Among the needs of adaptive solvers are greater repertoire of mesh entity adjacency relationships for h-refinement, parallel refinement and coarsening strategies as well as dynamic load balancing. The implementation of these routines need specialized data structures which provide fast updates not only on a single processor but also during entity migrations among processors. In addition, the portability requirement of the developed software is easily satisfied by the use of MPI message passing libraries.

So far the work has concentrated on improving and extending the earlier distributed grid library which was developed at Scientific Computation Research Center (SCOREC) of Rensselaer Polytechnic Institute. The extensions involved the implementation of a mesh database to which our distributed library was then linked. The dependence of entity migration procedures on the symmetric entity data structures, which provide two way links between entities of consecutive order (i.e. vertices, edges, faces and regions), was removed in order to save memory. 2D mesh refinement based on Rivara subdivision of elements was implemented. Unlike the Argonne National Lab's SUMAA3D Project implementation, our refinement routines do not require a solution to parallel graph coloring problem. In addition, the refinement routines were interfaced to the freely available IRIT solid modeller so that refined boundary edges could be snapped to actual geometry boundary. Extensive comparison between load balancing procedures based on incremental diffusive and global geometric coordinate bisection repartitioning have also been performed. It was found that the diffusive algorithms perform poorly in balancing refined meshes. This is mainly due to the high latency cost associated with the many steps taken to migrate the load and the inability to control the shape-quality of partitions by local decisions. Currently, 3D mesh refinement based on subdivision of elements is being implemented.

Future work will concentrate on demonstrating the use of the environment by coupling it with hp-based adaptive solvers.

ALEX POTHEN

Wavefront minimization of graphs and matrices

Combinatorial algorithms have been widely used in matrix computations in the past thirty years. However, algorithms from linear algebra have only rarely been used to solve combinatorial problems. Here we describe a combinatorial ordering problem called the wavefront minimization problem with applications in sparse matrix computations and computational biology, where a linear algebraic algorithm outperforms the best known combinatorial algorithm. Additionally the

algebraic algorithm is better able to cope when the input data is contaminated with errors, which is common in applications such as gene sequencing in biology.

In the wavefront minimization problem we are required to order the vertices of a graph (from 1 to n) such that at each step i the vertices numbered have as few unnumbered neighbors as possible. More precisely we seek to minimize, over all the steps, the sum of the number of unnumbered neighbors. We have described a hybrid algorithm that combines the spectral algorithm for wavefront reduction (Barnard, Pothen, and Simon (1995)) with the Sloan algorithm, which is currently the best combinatorial algorithm for reducing the wavefront. The hybrid algorithm further reduces the wavefronts over the spectral and Sloan algorithms. Our results show that these improved orderings lead to frontal solvers that are faster by an order of magnitude on advanced architectures such as the CRAY-YMP. One byproduct of our work is a more efficient implementation of the Sloan algorithm. We provide the first implementation of the Sloan algorithm whose running time is bounded by the number of nonzeros in the sparse matrix. The new implementation is an order of magnitude faster (from ten to forty fold) than earlier implementations of the Sloan algorithm on large problems. Our implementation of the Sloan algorithm required about two to three times the time needed by the simple RCM algorithm while delivering wavefronts that were about half of RCM on a collection of large test problems.

We are working on faster algorithms for wavefront minimization that do not sacrifice the quality of the spectral orderings, and on algorithms for weighted problems.

Symmetric indefinite systems of equations

Structural analysis applications give rise to large-scale linear systems of equations whose coefficient matrix is symmetric and indefinite. We have begun to design parallel algorithms and software for solving such systems of equations on parallel computers. This research is being conducted with partial support from IBM's University Partnership Program (UPP).

The issue here is that static data structures cannot be used in the numerical computation, since columns need to be permuted during the factorization depending on their numerical values. In the sparse context, it is possible that none of the available columns can be used as a pivot. We use block 1×1 and 2×2 pivots in our algorithm. We had to develop several new algorithmic ideas in this context to obtain a scalable parallel algorithm with fast execution times. We have completed the *first* parallel dense solver for this problem, and have nearly completed an initial implementation of a parallel sparse solver.

This project is at its initial stages; we expect to continue to develop an efficient scalable solver for this problem.

A microeconomic scheduler for parallel computers

We have considered the problem of scheduling on-line a set of jobs on a parallel computer with identical processors and have described an algorithm for this problem based on microeconomic ideas. This work will be useful at many parallel computer installations and workstation clusters in the world, where system administrators face the task of scheduling the jobs submitted on their machines. EASY, a job scheduler for parallel computers developed at Argonne National Labs by Lifka, and used at several institutions, uses the first-come-first-served policy with job reservations. This can be shown to be a special case of the microeconomic approach (when "income rates" are set to zero in this policy).

In recent work (with Ion Stoica) we have compared, through simulation experiments, the microeconomic scheduler that we have developed with other scheduling policies. We have designed a systematic set of experiments that explore various regions of the parameter space, and thereby to characterize robust scheduling policies. We explore three variants of the microeconomic approach that differ in the way income is distributed to a job. We show how they permit trade-offs between mutually antagonistic goals such as high system utilization and low user response times. The microeconomic approach has the additional advantages of maintaining fairness at the user level and providing each user with control over the performance of their jobs. Our work shows that the microeconomic approach is capable of modeling many commonly used scheduling policies. Furthermore, our results show that the microeconomic scheduler is more robust than the other schedulers included in this study.

We are continuing to refine the microeconomic scheduler for use in the parallel and cluster computing community, and to model its behavior analytically.

ARUN K. SOMANI

Specification methods and efficiency in analysis of regular interconnected systems

A performance and reliability tool must allow its users to transform the system description available to him/her from the architecture description easily. A system designer must be able to take his/her design documents and abstract it to a level suitable to analyze for performance and reliability. At the same time, the job of a performance and/or reliability analysis tool must be simplified by avoiding consideration of unnecessary states in the analysis. The system description methods, therefore, must include mechanisms to facilitate this. The goals of our research are to design efficient specification methods and languages to encourage engineers to use these tools for improving and verifying their design goals from the performance and reliability point of view.

Most designs are carried out in a hierarchical fashion. The user uses basic components to design subsystem which then in turn are used to design bigger subsystems and so on. The level of accuracy in performance and reliability modeling depends on how detailed model a user wishes to analyze. Hierarchical modeling methodology allows a user to specify and control the level of analysis relatively easily. The design process hierarchy can be directly used for the analysis as well. Our analysis methodology allows a user to specify the following efficiently: (1) component description, (2) system structure, (3) implied and derived behaviors, (4) symmetry and regularity in system interconnection structure. We exploit the symmetry in system interconnection structure in state space generation. With the symmetry, the number of states can be significantly reduced while fine behavior can be modeled at each transition if so desired.

We are currently developing tools and incorporating our techniques in them. A new tool called HIMAP is currently under development. Unique features of HIMAP provide immense flexibility to choose the modeling method appropriate to the system. HIMAP converts the high level representations to a Markov chain and uses existing tools to solve the Markov chain. We hope to finish this project by the end of year 1998.

MOULAY D. TIDRIRI

Schwarz-based methods in CFD

Our aim is the development and study of optimal two-level Schwarz/MUSCL type algorithms, which are intended for the parallel implicit solution of aerodynamics problems.

We have continued our investigation of Schwarz-based methods in CFD. This includes the combination of Schwarz algorithms with finite volume or finite element methods. Several practical and theoretical aspects of these combinations were addressed.

The refinement of both theoretical and numerical results we have obtained is underway.

LINDA F. WILSON

Automated load balancing in parallel discrete-event simulation

The performance of a parallel discrete-event simulation depends on the mapping of tasks to processors. In many cases, high performance is achieved only after rigorous fine-tuning is used to obtain an efficient mapping. In practice, good performance with minimal effort is often preferable to high performance with excessive effort.

In recent work, we modified the SPEEDES simulation framework to automate static load balancing. Using simulation models of queuing networks and the National Airspace System, we demonstrated that our automated load-balancing scheme can achieve good performance. In an effort to improve performance, we are currently examining various approaches for dynamic load balancing.

For future work, we will modify SPEEDES to implement dynamic load balancing. To accomplish this goal, we will also examine the related issue of object migration. This work is in collaboration with David Nicol.

REPORTS AND ABSTRACTS

Bataille, Françoise, Ye Zhou, and Jean-Pierre Bertoglio: *Energy transfer in compressible turbulence*. ICASE Report No. 95-65, October 4, 1995, 15 pages. Submitted to Physics of Fluids.

This *Letter* investigates the compressible energy transfer process. We extend a methodology developed originally for incompressible turbulence and use databases from numerical simulations of a weak compressible turbulence based on Eddy-Damped-Quasi-Normal-Markovian (EDQNM) closure. In order to analyze the compressible mode directly, the well known Helmholtz decomposition is used. While the compressible component has very little influence on the solenoidal part, we found that almost all of compressible turbulence energy is received from its solenoidal counterpart. We focus on the most fundamental building block of the energy transfer process, the triadic interactions. This analysis leads us to conclude that, at low turbulent Mach number, the compressible energy transfer process is dominated by a local radiative transfer (absorption) in both inertial and energy containing ranges.

Ristorcelli, J.R.: *Diagnostic statistics for the assessment and characterization of complex turbulent flows*. ICASE Report No. 95-67, December 13, 1995, 38 pages. Submitted to the Journal of Fluids Engineering.

A simple parameterization scheme for a complex turbulent flow using nondimensional parameters coming from the Reynolds stress equations is given. Definitions and brief descriptions of the physical significance of several nondimensional parameters that are used to characterize turbulence from the viewpoint of single-point turbulence closures are given. These nondimensional parameters reflect measures of 1) the spectral band width of the turbulence, 2) deviations from the ideal Kolmogorov behavior, 3) the relative magnitude, orientation, and temporal duration of the deformation to which the turbulence is subjected, 4) one and two-point measures of the large and small scale anisotropy of the turbulence and 5) inhomogeneity. This is in an attempt to create a more systematic methodology for the diagnosis and classification of turbulent flows as well as in the development, validation and application of turbulence model strategies. The parameters serve also to indicate the adequacy of various assumptions made in single-point turbulence models and in suggesting the appropriate turbulence strategy for a particular complex flow. The compilation will be of interest to experimentalists and to those involved in either computing turbulent flows or whose interests lie in verifying the adequacy of the phenomenological beliefs used in turbulence closures.

Kangro, Urve, and Roy Nicolaides: *Spurious fields in time domain computations of scattering problems*. ICASE Report No. 95-69, October 6, 1995, 15 pages. Submitted to IEEE Transactions on Antennas and Propagation.

In this paper two-dimensional electromagnetic scattering problems with a time-periodic incident field are considered. The scatterer is a perfect conductor, and an artificial boundary condition is used. The large time behavior of solutions, depending on (divergence-free) initial conditions, is characterized. It turns out that in addition to the expected time-periodic solution the limiting solution may also contain a spurious stationary field. The source of the stationary field is explained and equations describing it are obtained. Several avoidance strategies are discussed, and numerical comparisons of these techniques are given.

Hu, Fang Q.: *On absorbing boundary conditions for linear Euler equations by a perfectly matched layer*. ICASE Report No. 95-70, October 24, 1995, 28 pages. Submitted to the Journal of Computational Physics.

Recently, Berenger introduced a Perfectly Matched Layer (PML) technique for absorbing electromagnetic waves. In the present paper, a perfectly matched layer is proposed for absorbing out-going two-dimensional waves in a uniform mean flow, generated by linearized Euler equations. It is well known that the linearized Euler equations support acoustic waves, which travel with the speed of sound relative to the mean flow, and vorticity and entropy waves, which travel with the mean flow. The PML equations to be used at a region adjacent to the artificial boundary for absorbing these linear waves are defined. Plane wave solutions to the PML equations are developed and wave propagation and absorption properties are given. It is shown that the theoretical reflection coefficients at an interface between the Euler and PML domains are zero, independent of the angle of incidence and frequency of the waves. As such, the present study points out a possible alternative approach for absorbing out-going waves of the Euler equations with little or no reflection in computation. Numerical examples that demonstrate the validity of the proposed PML equations are also presented.

Ma, Kwan-Liu, Philip J. Smith, and Sandeep Jain: *Visualizing turbulent mixing of gases and particles*. ICASE Report No. 95-71, November 13, 1995, 15 pages. To appear in the Proceedings of the 7th International Symposium on Flow Visualization 1995.

A physical model and interactive computer graphics techniques have been developed for the visualization of the basic physical process of stochastic dispersion and mixing from steady-state CFD calculations. The mixing of massless particles and inertial particles is visualized by transforming the vector field from a traditionally Eulerian reference frame into a Lagrangian reference frame. Groups of particles are traced through the vector field for the mean path as well as their statistical dispersion about the mean position by using added scalar information about the root mean square value of the vector field and its Lagrangian time scale. In this way, clouds of particles in a turbulent environment

are traced, not just mean paths. In combustion simulations of many industrial processes, good mixing is required to achieve a sufficient degree of combustion efficiency. The ability to visualize this multiphase mixing can not only help identify poor mixing but also explain the mechanism for poor mixing. The information gained from the visualization can be used to improve the overall combustion efficiency in utility boilers or propulsion devices. We have used this technique to visualize steady-state simulations of the combustion performance in several furnace designs.

Mahalov, Alex, and Ye Zhou: *Analytical and phenomenological studies of rotating turbulence*. ICASE Report No. 95-72, October 6, 1995, 33 pages. Submitted to Physics of Fluids.

A framework, which combines mathematical analysis, closure theory, and phenomenological treatment, is developed to study the spectral transfer process and reduction of dimensionality in turbulent flows that are subject to rotation. First, we outline a mathematical procedure that is particularly appropriate for problems with two disparate time scales. The approach which is based on the Green's method leads to the Poincaré velocity variables and the Poincaré transformation when applied to rotating turbulence. The effects of the rotation are now reflected in the modifications to the convolution of nonlinear term. The Poincaré transformed equations are used to obtain a time-dependent analog of the Taylor-Proudman theorem valid in the asymptotic limit when the non-dimensional parameter $\mu \equiv \Omega t \rightarrow \infty$ (Ω is the rotation rate and t is the time). The 'split' of the energy transfer in both direct and inverse directions is established. Secondly, we apply the Eddy-Damped-Quasinormal-Markovian (EDQNM) closure to the Poincaré transformed Euler/Navier-Stokes equations. This closure leads to expressions for the spectral energy transfer. In particular, an unique triple velocity decorrelation time is derived with an explicit dependence on the rotation rate. This provides an important input for applying the phenomenological treatment of Zhou. In order to characterize the relative strength of rotation, another non-dimensional number, a spectral Rossby number, which is defined as the ratio of rotation and turbulence time scales, is introduced. Finally, the energy spectrum and the spectral eddy viscosity are deduced.

Jiang, Guang-Shan, and Chi-Wang Shu: *Efficient implementation of weighted ENO schemes*. ICASE Report No. 95-73, October 6, 1995, 41 pages. Submitted to the Journal of Computational Physics.

In this paper, we further analyze, test, modify and improve the high order WENO (weighted essentially non-oscillatory) finite difference schemes of Liu, Osher and Chan (1994). It was shown by Liu et al. that WENO schemes constructed from the r^{th} order (in L^1 norm) ENO schemes are $(r+1)^{th}$ order accurate. We propose a new way of measuring the smoothness of a numerical solution, emulating the idea of minimizing the total variation of the approximation, which results in a 5^{th} order WENO scheme for the case $r=3$, instead of the 4^{th} order with the original smoothness measurement by Liu et al. This 5^{th} order WENO scheme is as fast as the 4^{th} order WENO scheme of Liu et al. and, both schemes are about twice as fast as the 4^{th} order ENO schemes on vector supercomputers and as fast on serial and parallel computers. For Euler systems of gas dynamics,

we suggest to compute the weights from pressure and entropy instead of the characteristic values to simplify the costly characteristic procedure. The resulting WENO schemes are about twice as fast as the WENO schemes using the characteristic decompositions to compute weights, and work well for problems which do not contain strong shocks or strong reflected waves. We also prove that, for conservation laws with smooth solutions, all WENO schemes are convergent. Many numerical tests, including the 1D steady state nozzle flow problem and 2D shock entropy wave interaction problem, are presented to demonstrate the remarkable capability of the WENO schemes, especially the WENO scheme using the new smoothness measurement, in resolving complicated shock and flow structures. We have also applied Yang's artificial compression method to the WENO schemes to sharpen contact discontinuities.

Ma, Kwan-Liu: *Runtime volume visualization for parallel CFD*. ICASE Report No. 95-74, October 25, 1995, 13 pages. To appear in the Proceedings of the Parallel CFD '95 Conference.

This paper discusses some aspects of the design of a data distributed, massively parallel volume rendering library for runtime visualization of parallel computational fluid dynamics simulations in a message-passing environment. Unlike the traditional scheme in which visualization is a postprocessing step, the rendering is done in place on each node processor. Computational scientists who run large-scale simulations on a massively parallel computer can thus perform interactive monitoring of their simulations. The current library provides an interface to handle volume data on rectilinear grids. The same design principles can be generalized to handle other types of grids. For demonstration, we run a parallel Navier-Stokes solver making use of this rendering library on the Intel Paragon XP/S. The interactive visual response achieved is found to be very useful. Performance studies show that the parallel rendering process is scalable with the size of the simulation as well as with the parallel computer.

Ciardo, Gianfranco, Joshua Gluckman, and David Nicol: *Distributed state-space generation of discrete-state stochastic models*. ICASE Report No. 95-75, November 13, 1995, 25 pages. Submitted to the ORSA Journal on Computing.

High-level formalisms such as stochastic Petri nets can be used to model complex systems. Analysis of logical and numerical properties of these models often requires the generation and storage of the entire underlying state space. This imposes practical limitations on the types of systems which can be modeled. Because of the vast amount of memory consumed, we investigate distributed algorithms for the generation of state space graphs. The distributed construction allows us to take advantage of the combined memory readily available on a network of workstations. The key technical problem is to find effective methods for on-the-fly partitioning, so that the state space is evenly distributed among processors. In this paper we report on the implementation of a distributed state-space generator that may be linked to a number of existing system modeling tools. We discuss partitioning strategies in the context of Petri net models, and report on performance observed on a network of workstations, as well as on a distributed memory multi-computer.

Lewis, Robert Michael: *Practical aspects of variable reduction formulations and reduced basis algorithms in multidisciplinary design optimization*. ICASE Report No. 95-76, November 20, 1995, 19 pages. To appear in the Proceedings of the ICASE/LaRC Workshop on Multidisciplinary Design Optimization.

This paper discusses certain connections between nonlinear programming algorithms and the formulation of optimization problems for systems governed by state constraints. I work through the calculation of the sensitivities associated with the different formulations and present some useful relationships between them. These relationships have practical consequences; if one uses a reduced basis nonlinear programming algorithm, then the implementations for the different formulations need only differ in a single step.

Zhou, Ye, and Robert Rubinstein: *Sweeping and straining effects in sound generation by high Reynolds number isotropic turbulence*. ICASE Report No. 95-77, November 20, 1995, 11 pages. Submitted to Physics of Fluids.

The sound radiated by isotropic turbulence is computed using inertial range scaling expressions for the relevant two time and two point correlations. The result depends on whether the decay of Eulerian time correlations is dominated by large scale sweeping or by local straining: the straining hypothesis leads to an expression for total acoustic power given originally by Proudman, whereas the sweeping hypothesis leads to a more recent result due to Lilley.

Somani, Arun K., and Tianming Zhang: *Architecture and performance analysis of DIRSMIN: A fault-tolerant switch using dilated reduced-stage MIN*. ICASE Report No. 95-78, November 20, 1995, 29 pages. To be submitted to IEEE Transactions on Communications.

We develop and analyze a dilated high performance fault tolerant fast packet multistage interconnection network (MIN) in this paper. In this new design, the links at the input and the output stages of a dilated banyan-based MIN are rearranged to create multiple routes for each source-destination pair in the network after removing one stage in the network. These multiple paths are link- and node-disjoint. Fault tolerance at low latency is achieved by sending multiple copies of each input packet simultaneously using different routes and different priorities. This guarantees that high throughput is maintained even in the presence of faults. Throughput is analyzed using simulation and analysis and we show that the new design has considerably higher performance in the presence of a faulty switching element (SE) or link in comparison to dilated networks. We also analyze the reliability and show that the new design has superior reliability in comparison to competing proposals.

D'Ambrosio, Domenic, and Roberto Marsilio: *Shock-induced separated structures in symmetric corner flows*. ICASE Report No. 95-79, November 20, 1995, 32 pages. To be submitted to the Journal of Fluid Mechanics.

Three-dimensional supersonic viscous laminar flows over symmetric corners are considered in this paper. The characteristic features of such configurations are discussed and an historical survey on the past research work is presented. A new contribution based on a numerical technique that solves the parabolized form of the Navier-Stokes equations is presented. Such a method makes it possible to obtain very detailed descriptions of the flowfield with relatively modest CPU time and memory storage requirements. The numerical approach is based on a space-marching technique, uses a finite volume discretization and an upwind flux-difference splitting scheme (developed for the steady flow equations) for the evaluation of the inviscid fluxes. Second order accuracy is reached following the guidelines of the ENO schemes. Different free-stream conditions and geometrical configurations are considered. Primary and secondary streamwise vortical structures embedded in the boundary layer and originated by the interaction of the latter with shock waves are detected and studied. Computed results are compared with experimental data taken from literature.

Wilson, Linda F., and David M. Nicol: *Experiments in automated load balancing*. ICASE Report No. 95-80, December 6, 1995, 15 pages. Submitted to PADS '96: 10th Workshop on Parallel and Distributed Simulation.

One of the promises of parallelized discrete-event simulation is that it might provide significant speedups over sequential simulation. In reality, high performance cannot be achieved unless the system is fine-tuned to balance computation, communication, and synchronization requirements. As a result, parallel discrete-event simulation needs tools to automate the tuning process with little or no modification to the user's simulation code.

In this paper, we discuss our experiments in automated load balancing using the SPEEDES simulation framework. Specifically, we examine three mapping algorithms that use run-time measurements. Using simulation models of queuing networks and the National Airspace System, we investigate (i) the use of run-time data to guide mapping, (ii) the utility of considering communication costs in a mapping algorithm, (iii) the degree to which computational "hot-spots" ought to be broken up in the linearization, and (iv) the relative execution costs of the different algorithms.

Dickens, Phillip, Matthew Haines, Piyush Mehrotra, and David Nicol: *Towards a thread-based parallel direct execution simulator*. ICASE Report No. 95-81, December 4, 1995, 17 pages. To appear in the 29th Hawaii International Conference on Computer Systems.

Parallel direct execution simulation is an important tool for performance and scalability analysis of large message passing parallel programs executing on top of a *virtual computer*. However, detailed simulation of message-passing codes requires a great deal of computation. We are therefore interested in pursuing implementation techniques which can decrease this cost. One idea is

to implement the application virtual processes as lightweight threads rather than traditional Unix processes, reducing both on-processor communication costs and context-switching costs. In this paper we describe an initial implementation of a thread-based parallel direct execution simulator. We discuss the advantages of such an approach and present preliminary results that indicate a significant improvement over the process-based approach.

Girimaji, Sharath S.: *Fully-explicit and self-consistent algebraic Reynolds stress model*. ICASE Report No. 95-82, December 31, 1995, 24 pages. Submitted to the Journal of Fluid Mechanics.

A fully-explicit, self-consistent algebraic expression for the Reynolds stress, which is the exact solution to the Reynolds stress transport equation in the ‘weak equilibrium’ limit for two-dimensional mean flows for all linear and some quasi-linear pressure-strain models, is derived. Current explicit algebraic Reynolds stress models derived by employing the ‘weak equilibrium’ assumption treat the production-to-dissipation (P/ε) ratio implicitly, resulting in an effective viscosity that can be singular away from the equilibrium limit. In the present paper, the set of simultaneous algebraic Reynolds stress equations are solved in the full non-linear form and the eddy viscosity is found to be non-singular. Preliminary tests indicate that the model performs adequately, even for three dimensional mean flow cases. Due to the explicit and non-singular nature of the effective viscosity, this model should mitigate many of the difficulties encountered in computing complex turbulent flows with the algebraic Reynolds stress models.

Chrisochoides, Nikos: *Multithreaded model for dynamic load balancing parallel adaptive PDE computations*. ICASE Report No. 95-83, February 29, 1996, 25 pages. To appear in the Applied Numerical Mathematics Journal.

We present a multithreaded model for the dynamic load-balancing of numerical, adaptive computations required for the solution of Partial Differential Equations (PDEs) on multiprocessors. Multithreading is used as a means of exploring concurrency at the processor level in order to tolerate synchronization costs inherent to traditional (non-threaded) parallel adaptive PDE solvers. Our preliminary analysis for parallel, adaptive PDE solvers indicates that multithreading can be used as a mechanism to mask overheads required for the dynamic balancing of processor workloads with computations required for the actual numerical solution of the PDEs. Also, multithreading can simplify the implementation of dynamic load-balancing algorithms, a task that is very difficult for traditional data parallel adaptive PDE computations. Unfortunately, multithreading does not always simplify program complexity, often makes code re-usability difficult, and increases software complexity.

Arian, Eyal: *Analysis of the Hessian for aeroelastic optimization*. ICASE Report No. 95-84, December 31, 1995, 22 pages. To be submitted to the Journal of Computational Physics.

The symbol of the Hessian for an aeroelastic optimization model problem is analyzed for the optimization of a plate's shape and rigidity distribution with respect to a given cost function. The flow is modeled by the small-disturbance full-potential equation and the structure is modeled by an isotropic (von Kármán) plate equation. The cost function consists of both aerodynamic and structural terms. In the new analysis the symbol of the cost function Hessian near the minimum is computed. The result indicates that under some conditions, which are likely fulfilled in most applications, the system is decoupled for the non-smooth components. The result also shows that the structure part in the Hessian is well-conditioned while the aerodynamic part is ill-conditioned. Applications of the result to optimization strategies are discussed.

Nicol, David M., Daniel L. Palumbo, and Michael Ulrey: *Integrating reliability analysis with a performance tool*. ICASE Report No. 95-85, December 31, 1995, 19 pages. Submitted to Communications on Reliability, Maintainability and Supportability.

A large number of commercial simulation tools support performance oriented studies of complex computer and communication systems. Reliability of these systems, when desired, must be obtained by remodeling the system in a different tool. This has obvious drawbacks: (i) substantial extra effort is required to create the reliability model, (ii) through modeling error the reliability model may not reflect precisely the same system as the performance model, (iii) as the performance model evolves one must continuously reevaluate the validity of assumptions made in that model. In this paper we describe an approach—and a tool that implements this approach—for integrating a reliability analysis engine into a production quality simulation based performance modeling tool, and for modeling within such an integrated tool. The integrated tool allows one to use the same modeling formalisms to conduct both performance and reliability studies. We describe how the reliability analysis engine is integrated into the performance tool, describe the extensions made to the performance tool to support the reliability analysis, and consider the tool's performance.

Cai, Xiao-Chuan, and Marcus Sarkis: *Local multiplicative Schwarz algorithm for convection-diffusion equations*. ICASE Report No. 95-86, December 31, 1995, 15 pages. Submitted to SIAM Journal of Scientific Computing.

We develop a new class of overlapping Schwarz type algorithms for solving scalar convection-diffusion equations discretized by finite element or finite difference methods. The preconditioners consist of two components, namely, the usual two-level additive Schwarz preconditioner and the sum of some quadratic terms constructed by using products of ordered neighboring subdomain preconditioners. The ordering of the subdomain preconditioners is determined by considering the direction of the flow. We prove that the algorithms are optimal in the sense that the convergence rates are independent of the mesh size, as well as the number of subdomains. We show by numerical

examples that the new algorithms are less sensitive to the direction of the flow than either the classical multiplicative Schwarz algorithms, and converge faster than the additive Schwarz algorithms. Thus, the new algorithms are more suitable for fluid flow applications than the classical additive or multiplicative Schwarz algorithms.

Cai, Xiao-Chuan, David E. Keyes, and V. Venkatakrishnan: *Newton-Krylov-Schwarz: An implicit solver for CFD*. ICASE Report No. 95-87, December 31, 1995, 19 pages. Submitted to the Proceedings of the 8th International Conference on Domain Decomposition Methods.

Newton-Krylov methods and Krylov-Schwarz (domain decomposition) methods have begun to become established in computational fluid dynamics (CFD) over the past decade. The former employ a Krylov method inside of Newton's method in a Jacobian-free manner, through directional differencing. The latter employ an overlapping Schwarz domain decomposition to derive a preconditioner for the Krylov accelerator that relies primarily on local information, for data-parallel concurrency. They may be composed as Newton-Krylov-Schwarz (NKS) methods, which seem particularly well suited for solving nonlinear elliptic systems in high-latency, distributed-memory environments. We give a brief description of this family of algorithms, with an emphasis on domain decomposition iterative aspects. We then describe numerical simulations with Newton-Krylov-Schwarz methods on aerodynamics applications emphasizing comparisons with a standard defect-correction approach, subdomain preconditioner consistency, subdomain preconditioner quality, and the effect of a coarse grid.

Wilson, Linda F.: *Analysis of algorithmic structures with heterogeneous tasks*. ICASE Report No. 96-1, (NASA CR-198259), January 29, 1996, 23 pages. To be presented at the International Phoenix Conference on Computers and Communications.

Developing efficient programs for distributed systems is difficult because computations must be efficiently distributed and managed on multiple processors. In particular, the programmer must partition functions and data in an attempt to find a reasonable balance between parallelism and overhead. Furthermore, it is very expensive to code an algorithm only to find out that the implementation is not efficient. As a result, it is often necessary to determine and examine those characteristics of an algorithm that can be used to predict its suitability for a distributed computing system.

In earlier work (Wilson, 1994 and Wilson and Gonzalez, 1994), we presented a framework for the study of synchronization and communication effects on the theoretical performance of common homogeneous algorithmic structures. In particular, we examined the synchronous, asynchronous, nearest-neighbor, and asynchronous master-slave structures in terms of expected execution times. In this paper, we examine the effects of synchronization and communication on the expected execution times of heterogeneous algorithmic structures. Specifically, we consider structures containing two different types of tasks, where the execution times of the tasks follow one of two different uniform distributions or one of two different normal distributions. Furthermore, we compare the

expected execution times of the heterogeneous algorithmic structures with times for corresponding homogeneous structures. Finally, we develop bounds for the expected execution times of the heterogeneous structures and compare those bounds to simulated execution times.

Rubinstein, Robert: *Inertial range dynamics in Boussinesq turbulence*. ICASE Report No. 96-2, (NASA CR-198260), February 9, 1996, 22 pages. To be submitted to Physics of Fluids.

L'vov and Falkovich (Physica D 57) have shown that the dimensionally possible inertial range scaling laws for Boussinesq turbulence, Kolmogorov and Bolgiano scaling, describe steady states with, respectively, constant flux of kinetic energy and of entropy. Following Woodruff (Phys. Fluids 6), these scaling laws are treated as similarity solutions of the direct interaction approximation for Boussinesq turbulence. The Kolmogorov scaling solution corresponds to a weak perturbation by gravity of a state in which the temperature is a passive scalar but in which a source of temperature fluctuations exists. Using standard inertial range balances, the effective viscosity and conductivity, turbulent Prandtl number, and spectral scaling law constants are computed for Bolgiano scaling.

Chiavassa, G., and J. Liandrat: *On the effective construction of compactly supported wavelets satisfying homogeneous boundary conditions on the interval*. ICASE Report No. 96-3, (NASA CR-198261), January 29, 1996, 33 pages. Submitted to Applied and Computational Harmonic Analysis.

We construct compactly supported wavelet bases satisfying homogeneous boundary conditions on the interval $[0,1]$. The maximum features of multiresolution analysis on the line are retained, including polynomial approximation and tree algorithms. The case of $H_0^1([0,1])$ is detailed, and numerical values, required for the implementation, are provided for the Neumann and Dirichlet boundary conditions.

Tidriri, M.D.: *Schwarz-based algorithms for compressible flows*. ICASE Report No. 96-4, (NASA CR-198273), February 9, 1996, 38 pages. To be submitted to Computer Methods in Applied Mechanics and Engineering.

We investigate in this paper the application of Schwarz-based algorithms to compressible flows. First, we study the combination of these methods with defect-correction procedures. We then study the effect on the Schwarz-based methods of replacing the explicit treatment of the boundary conditions by an implicit one. In the last part of this paper we study the combination of these methods with Newton-Krylov matrix-free methods. Numerical experiments that show the performance of our approaches are then presented.

Horton, Graham, Vidyadhar G. Kulkarni, David M. Nicol, and Kishor S. Trivedi: *Fluid stochastic Petri nets: Theory, applications, and solution*. ICASE Report No. 96-5, (NASA CR-198274), February 9, 1996, 21 pages. Submitted to the European Journal on Operations Research.

In this paper we introduce a new class of stochastic Petri nets in which one or more places can hold fluid rather than discrete tokens. We define a class of fluid stochastic Petri nets in such a way that the discrete and continuous portions may affect each other. Following this definition we provide equations for their transient and steady-state behavior. We present several examples showing the utility of the construct in communication network modeling and reliability analysis, and discuss important special cases. We then discuss numerical methods for computing the transient behavior of such nets. Finally, some numerical examples are presented.

Pavarino, Luca F.: *Domain decomposition algorithms for first-order system least squares methods*. ICASE Report No. 96-6, (NASA CR-198275), February 21, 1996, 19 pages. Submitted to the SIAM Journal of Scientific Computing.

Least squares methods based on first-order systems have been recently proposed and analyzed for second-order elliptic equations and systems. They produce symmetric and positive definite discrete systems by using standard finite element spaces, which are not required to satisfy the inf-sup condition. In this paper, several domain decomposition algorithms for these first-order least squares methods are studied. Some representative overlapping and substructuring algorithms are considered in their additive and multiplicative variants. The theoretical and numerical results obtained show that the classical convergence bounds (on the iteration operator) for standard Galerkin discretizations are also valid for least squares methods.

Dubois, Thierry, Francois Jauberteau, and Ye Zhou: *Influences of subgrid scale dynamics on resolvable scale statistics in large-eddy simulations*. ICASE Report No. 96-7, (NASA CR-198278), February 21, 1996, 28 pages. Submitted to Physica D.

Recently, the ϵ -expansion and recursive renormalization group (RNG) theories as well as approximation inertial manifolds (AIM) have been exploited as means of systematically modeling subgrid scales in large-eddy simulations (LES). Although these theoretical approaches are rather complicated mathematically, their key approximations can be investigated using direct numerical simulations (DNS). In fact, the differences among these theories can be traced to whether they retain or neglect interactions between the subgrid-subgrid and subgrid-resolvable scales. In this paper, we focus on the influence of these two interactions on the evolution of the resolvable scales in LES: the effect^A which keeps only the interactions between the small and large scales; and, the effect^B which, on the other hand, keeps only the interactions among the subgrid-subgrid scales. The performance of these models is analyzed using the velocity fields of the direct numerical simulations. Specifically, our comparison is based on the analysis of the energy and enstrophy spectra, as well as higher-order statistics of the velocity and velocity derivatives. We found that the energy

spectrum and higher-order statistics for the simulations with the effect^A (referred to, hereafter, as model^A) are in very good agreement with the filtered DNS. The comparison between the computations with effect^B (referred to, hereafter, as model^B) and the filtered DNS, however, is not satisfactory. Moreover, the decorrelation between the filtered DNS and model^A is much slower than that of the filtered DNS and model^B. Therefore, we conclude that the model^A, taking into account the interactions between the subgrid and resolvable scales, is a faithful subgrid model for LES for the range of Reynolds numbers considered.

Abarbanel, Saul, and Adi Ditkowski: *Multi-dimensional asymptotically stable 4th-order accurate schemes for the diffusion equation*. ICASE Report No. 96-8, (NASA CR-198279), February 29, 1996, 37 pages. To be submitted to the Journal of Computational Physics.

An algorithm is presented which solves the multi-dimensional diffusion equation on complex shapes to the 4th-order accuracy and is asymptotically stable in time. This bounded-error result is achieved by constructing, on a rectangular grid, a differentiation matrix whose symmetric part is negative definite. The differentiation matrix accounts for the Dirichlet boundary condition by imposing penalty like terms.

Numerical examples in 2D show that the method is effective even where standard schemes, stable by traditional definitions, fail.

Hall, Philip, and Trudi A. Shortis: *On the effect of feedback control on Bénard convection in a Boussinesq fluid*. ICASE Report No. 96-9, (NASA CR-198280), March 18, 1996, 26 pages. Submitted to Phil. Trans. Roy. Soc.

The effect of nonlinear feedback control strategies on the platform of convection in a Boussinesq fluid heated from below is investigated. In the absence of the control, given that non-Boussinesq effects may be neglected, it is well known that convection begins in the form of a supercritical bifurcation to rolls. Non-Boussinesq behaviour destroys the symmetry of the basic state, and through a subcritical bifurcation leads to the formation of hexagonal cells. Here we discuss the influence of regulation of the lower surface temperature by means of a control mechanism, made up of a combination of a proportional linear and nonlinear controller, on the stability of the hexagonal cell pattern.

Hewitt, R.E., and P. Hall: *The evolution of finite amplitude wavetrains in plane channel flow*. ICASE Report No. 96-10, (NASA CR-198281), February 26, 1996, 35 pages. Submitted to Phil. Trans. Roy. Soc.

We consider a viscous incompressible fluid flow driven between two parallel plates by a constant pressure gradient. The flow is at a finite Reynolds number, with an $O(1)$ disturbance in the form of a traveling wave. A phase equation approach is used to discuss the evolution of slowly varying fully

nonlinear two dimensional wavetrains. We consider uniform wavetrains in detail, showing that the development of a wavenumber perturbation is governed by Burgers equation in most cases. The wavenumber perturbation theory, constructed using the phase equation approach for a uniform wavetrain, is shown to be distinct from an amplitude perturbation expansion about the periodic flow. In fact we show that the amplitude equation contains only linear terms and is simply the heat equation. We review, briefly, the well known dynamics of Burgers equation, which imply that both shock structures and finite time singularities of the wavenumber perturbation can occur with respect to the slow scales. Numerical computations have been performed to identify areas of the {wavenumber, Reynolds number, energy} neutral surface for which each of these possibilities can occur. We note that the evolution equations will breakdown under certain circumstances, in particular for a weakly nonlinear secondary flow. Finally we extend the theory to three dimensions and discuss the limit of a weak spanwise dependence for uniform wavetrains, showing that two functions are required to describe the evolution. These unknowns are a phase and a pressure function which satisfy a pair of linearly coupled partial differential equations. The results obtained from applying the same analysis to the fully three dimensional problem are included as an appendix.

Fenno, C.C., Jr., A. Bayliss, and L. Maestrello: *Response of a panel structure forced by the noise from a nearly sonic jet*. ICASE Report No. 96-11, (NASA CR-198282), February 29, 1996, 22 pages. Submitted to the AIAA Journal and to appear at the AIAA Aeroacoustics Conference, March 1996.

A model of a high subsonic jet with a nearby array of flexible, aircraft-type panels is studied numerically in two dimensions. The jet is excited by a limited duration, spatially localized starter pulse in the potential core. The long time evolution of unsteady disturbances in the jet, the responses of the panels and the ensuing radiation are computed. The results show that the spectral response of both the jet and the panels is concentrated in a relatively narrow frequency band centered at a Strouhal number (based on jet exit velocity) of approximately 0.25 and associated harmonics. The loading on the panels generally increases with downstream distance. Panel radiation is weakest in upstream directions. Interior zones of silence, due to destructive interference of radiation from the panels, are observed.

Smyrlis, Yiorgos S., and Demetrios T. Papageorgiou: *Computational study of chaotic and ordered solutions of the Kuramoto-Sivashinsky equation*. ICASE Report No. 96-12, (NASA CR-198283), February 21, 1996, 34 pages. To appear in Advances in Multi-Fluid Flows. Proceedings of the AMS-IMS-SIAM Joint Summer Research Conference on Multi-Fluid Flows and Interfacial Instabilities.

We report the results of extensive numerical experiments on the Kuramoto-Sivashinsky equation in the strongly chaotic regime as the viscosity parameter is decreased and increasingly more linearly unstable modes enter the dynamics. General initial conditions are used and evolving states do not assume odd-parity. A large number of numerical experiments are employed in order to obtain quantitative characteristics of the dynamics. We report on different routes to chaos and provide

numerical evidence and construction of strange attractors with self-similar characteristics. As the "viscosity" parameter decreases the dynamics becomes increasingly more complicated and chaotic. In particular it is found that regular behavior in the form of steady state or steady state traveling waves is supported amidst the time-dependent and irregular motions. We show that multimodal steady states emerge and are supported on decreasing windows in parameter space. In addition we invoke a self-similarity property of the equation, to show that these profiles are obtainable from global fixed point attractors of the Kuramoto-Sivashinsky equation at much larger values of the viscosity.

Ponenti, Pj., and J. Liandrat: *Numerical algorithms based on biorthogonal wavelets*. ICASE Report No. 96-13, (NASA CR-198290), March 4, 1996, 35 pages. Submitted to *Numerische Mathematik*.

Wavelet bases are used to generate spaces of approximation for the resolution of bidimensional elliptic and parabolic problems. Under some specific hypotheses relating the properties of the wavelets to the order of the involved operators, it is shown that an approximate solution can be built. This approximation is then stable and converges towards the exact solution. It is designed such that fast algorithms involving biorthogonal multi resolution analyses can be used to resolve the corresponding numerical problems.

Detailed algorithms are provided as well as the results of numerical tests on partial differential equations defined on the bidimensional torus.

Rubinstein, Robert, and Ye Zhou: *Analytical theory of the destruction terms in dissipation rate transport equations*. ICASE Report No. 96-14, (NASA CR-198291), March 13, 1996, 18 pages. Submitted to *Physics of Fluids*.

Modeled dissipation rate transport equations are often derived by invoking various hypotheses to close correlations in the corresponding exact equations. D.C. Leslie suggested that these models might be derived instead from Kraichnan's wavenumber space integrals for inertial range transport power. This suggestion is applied to the destruction terms in the dissipation rate equations for incompressible turbulence, buoyant turbulence, rotating incompressible turbulence, and rotating buoyant turbulence. Model constants like $C_{\epsilon 2}$ are expressed as integrals; convergence of these integrals implies the absence of Reynolds number dependence in the corresponding destruction term. The dependence of $C_{\epsilon 2}$ on rotation rate emerges naturally; sensitization of the modeled dissipation rate equation to rotation is not required. A buoyancy related effect which is absent in the exact transport equation for temperature variance dissipation, but which sometimes improves computational predictions, also arises naturally. Both the presence of this effect and the appropriate time scale in the modeled transport equation depend on whether Bolgiano or Kolmogorov inertial range scaling applies. A simple application of these methods leads to a preliminary dissipation rate equation for rotating buoyant turbulence.

Kopriva, David A.: *A conservative staggered-grid Chebyshev multidomain method for compressible flows. II: A semi-structured method.* ICASE Report No. 96-15, (NASA CR-198292), March 14, 1996, 27 pages. To be submitted to the Journal of Computational Physics.

We present a Chebyshev multidomain method that can solve systems of hyperbolic equations in conservation form on an unrestricted quadrilateral subdivision of a domain. Within each subdomain the solutions and fluxes are approximated by a staggered-grid Chebyshev method. Thus, the method is unstructured within subdomains. Communication between subdomains is done by a mortar method in such a way that the method is globally conservative. The method is applied to both linear and non-linear test problems and spectral accuracy is demonstrated.

Newman, Perry A., Gene J.-W. Hou, and Arthur C. Taylor, III: *Observations regarding use of advanced CFD analysis, sensitivity analysis, and design codes in MDO.* ICASE Report No. 96-16, (NASA CR-198293), March 14, 1996, 18 pages. Submitted to SIAM for Proceedings of the Multidisciplinary Design Optimization: State of the Art.

Observations regarding the use of advanced computational fluid dynamics (CFD) analysis, sensitivity analysis (SA), and design codes in gradient-based multidisciplinary design optimization (MDO) reflect our perception of the interactions required of CFD and our experience in recent aerodynamic design optimization studies using CFD. Sample results from these latter studies are summarized for conventional optimization (analysis-SA codes) and simultaneous analysis and design optimization (design code) using both Euler and Navier-Stokes flow approximations. The SA codes is greater than that required for design codes. Thus, an MDO formulation that utilizes the more efficient design codes where possible is desired. However, in the aerovehicle MDO problem, the various disciplines that are involved have different design points in the flight envelope; therefore, CFD analysis-SA codes are required at the aerodynamic "off design" points. The suggested MDO formulation is a hybrid multilevel optimization procedure that consists of both multipoint CFD analysis-SA codes and multipoint CFD design codes that perform suboptimizations.

Papageorgiou, Demetrios T.: *Description of jet breakup.* ICASE Report No. 96-17, (NASA CR-198294), March 14, 1996, 39 pages. To appear in Advances in Multi-Fluid Flows. Proceedings of the AMS-IMS-SIAM Joint Summer Research Conference on Multi-Fluid Flows and Interfacial Instabilities.

In this article we review recent results on the breakup of cylindrical jets of a Newtonian fluid. Capillary forces provide the main driving mechanism and our interest is in the description of the flow as the jet pinches to form drops. The approach is to describe such topological singularities by constructing local (in time and space) similarity solutions from the governing equations. This is described for breakup according to the Euler, Stokes or Navier-Stokes equations. It is found that slender jet theories can be applied when viscosity is present, but for inviscid jets the local shape of the jet at breakup is most likely of a non-slender geometry. Systems of one-dimensional models of the governing equations are solved numerically in order to illustrate these differences.

Cziesla, T., H. Braun, G. Biswas, and N.K. Mitra: *Large eddy simulation in a channel with exit boundary conditions*. ICASE Report No. 96-18, (NASA CR-198304), March 18, 1996, 21 pages. To be submitted to the Journal of Fluids Engineering.

The influence of the exit boundary conditions (vanishing first derivative of the velocity components and constant pressure) on the large eddy simulation of the fully developed turbulent channel flow has been investigated for equidistant and stretched grids at the channel exit.

Results show that the chosen exit boundary conditions introduce some small disturbance which is mostly damped by the grid stretching. The difference between the fully developed turbulent channel flow obtained with LES with periodicity condition and the inlet and exit and the LES with fully developed flow at the inlet and the exit boundary condition is less than 10% for equidistant grids and less than 5% for the case grid stretching. The chosen boundary condition is of interest because it may be used in complex flows with backflow at exit.

Ristorcelli, J.R., and J.H. Morrison: *The Favre-Reynolds average distinction and a consistent gradient transport expression for the dissipation*. ICASE Report No. 96-19, (NASA CR-198305), March 20, 1996, 7 pages. Submitted to Physics of Fluids.

Two equation and higher order closures for compressible turbulence fail to capture the compressible wall layers' log scaling. Accounting for the distinction between Favre and Reynolds averaged variables in the compressible moment equations indicate that turbulent transport expressions obtained using the "variable density approximation" are in error. The error is related to the enstrophy, a *Reynolds averaged* variable appearing in the equation for the *Favre averaged* k ; recognizing this fact an expression for the transport of dissipation consistent with simple mixing length arguments is obtained. Within the (limited) context of a gradient transport hypothesis a rational form for the turbulent transport of the dissipation is found. Modestly better agreement with the well established compressible Van Driest log scaling is found in $k - \varepsilon$ calculation.

Lewis, Robert Michael, and Virginia Torczon: *Pattern search algorithms for bound constrained minimization*. ICASE Report No. 96-20, (NASA CR-198306), March 20, 1996, 21 pages. Submitted to the SIAM Journal on Optimization.

We present a convergence theory for pattern search methods for solving bound constrained nonlinear programs. The analysis relies on the abstract structure of pattern search methods and an understanding of how the pattern interacts with the bound constraints. This analysis makes it possible to develop pattern search methods for bound constrained problems while only slightly restricting the flexibility present in pattern search methods for unconstrained problems. We prove global convergence despite the fact that pattern search methods do not have explicit information concerning the gradient and its projection onto the feasible region and consequently are unable to enforce explicitly a notion of sufficient feasible decrease.

Wu, Jie-Zhi, Ye Zhou, and Jian-Ming Wu: *Reduced stress tensor and dissipation and the transport of Lamb vector*. ICASE Report No. 96-21, (NASA CR-198307), March 20, 1996, 13 pages. Submitted to Physics of Fluids.

We develop a methodology to ensure that the stress tensor, regardless of its number of independent components, can be reduced to an exactly equivalent one which has the same number of independent components as the surface force. It is applicable to the momentum balance if the shear viscosity is constant. A direct application of this method to the energy balance also leads to a reduction of the dissipation rate of kinetic energy. Following this procedure, significant saving in analysis and computation may be achieved. For turbulent flows, this strategy immediately implies that a given Reynolds stress model can always be replaced by a reduced one before putting it into computation. Furthermore, we show how the modeling of Reynolds stress tensor can be reduced to that of the mean turbulent Lamb vector alone, which is much simpler. As a first step of this alternative modeling development, we derive the governing equations for the Lamb vector and its square. These equations form a basis of new second-order closure schemes and, we believe, should be favorably compared to that of traditional Reynolds stress transport equation.

Jayasimha, D.N., M.E. Hayder, and S.K. Pillay: *An evaluation of architectural platforms for parallel Navier-Stokes computations*. ICASE Report No. 96-22, (NASA CR-198308), March 26, 1996, 22 pages. Submitted to the Journal of Supercomputing.

We study the computational, communication, and scalability characteristics of a Computational Fluid Dynamics application, which solves the time accurate flow field of a jet using the compressible Navier-Stokes equations, on a variety of parallel architecture platforms. The platforms chosen for this study are a cluster of workstations (the LACE experimental testbed at NASA Lewis), a shared memory multiprocessor (the Cray YMP), and distributed memory multiprocessors with different topologies — the IBM SP and the Cray T3D. We investigate the impact of various networks connecting the cluster of workstations on the performance of the application and the overheads induced by popular message passing libraries used for parallelization. The work also highlights the importance of matching the memory bandwidth to the processor speed for good single processor performance. By studying the performance of an application on a variety of architectures, we are able to point out the strengths and weaknesses of each of the example computing platforms.

Somani, Arun K.: *Simplified phased-mission system analysis for systems with independent component repairs*. ICASE Report No. 96-23, (NASA CR-198318), March 27, 1996, 25 pages. To be submitted to Sigmetrics 1996.

Accurate analysis of reliability of system requires that it accounts for all major variations in system's operation. Most reliability analyses assume that the system configuration, success criteria, and component behavior remain the same. However, multiple phases are natural. We present a new computationally efficient technique for analysis of phased-mission systems where the operational

states of a system can be described by combinations of components states (such as fault trees or assertions). Moreover, individual components may be repaired, if failed, as part of system operation but repairs are independent of the system state. For repairable systems Markov analysis techniques are used but they suffer from state space explosion. That limits the size of system that can be analyzed and it is expensive in computation. We avoid the state space explosion. The phase algebra is used to account for the effects of variable configurations, repairs, and success criteria from phase to phase. Our technique yields exact (as opposed to approximate) results. We demonstrate our technique by means of several examples and present numerical results to show the effects of phases and repairs on the system reliability/availability.

Shebalin, John V., and Stephen L. Woodruff: *Kolmogorov flow in three dimensions*. ICASE Report No. 96-24, (NASA CR-198319), March 20, 1996, 25 pages. To be submitted to Physics of Fluids.

A numerical study of the long-time evolution of incompressible Navier-Stokes turbulence forced at a single long-wavelength Fourier mode, *i.e.*, a Kolmogorov flow, has been completed. The boundary conditions are periodic in three dimensions and the forcing is effected by imposing a steady, two-dimensional, sinusoidal shear velocity which is directed along the x -direction and varies along the z -direction. A comparison with experimental data shows agreement with measured cross-correlations of the turbulent velocity components which lie in the mean-flow plane. A statistical analysis reveals that the shear-driven turbulence studied here has significant spectral anisotropy which increases with wave number.

del Rosario, R.C.H., and R.C. Smith: *Spline approximation of thin shell dynamics*. ICASE Report No. 96-26, (NASA CR-198323), March 27, 1996, 39 pages. Submitted to International Journal for Numerical Methods in Engineering.

A spline-based method for approximating thin shell dynamics is presented here. While the method is developed in the context of the Donnell-Mushtari thin shell equations, it can be easily extended to the Byrne-Flügge-Lur'ye equations or other models for shells of revolution as warranted by applications. The primary requirements for the method include accuracy, flexibility and efficiency in smart material applications. To accomplish this, the method was designed to be flexible with regard to boundary conditions, material nonhomogeneities due to sensors and actuators, and inputs from smart material actuators such as piezoceramic patches. The accuracy of the method was also of primary concern, both to guarantee full resolution of structural dynamics and to facilitate the development of PDE-based controllers which ultimately require real-time implementation. Several numerical examples provide initial evidence demonstrating the efficacy of the method.

INTERIM REPORTS

Bokhari, Shahid H.: *Communication overhead on the Intel Paragon, IBM SP2 & Meiko CS-2.*
ICASE Interim Report No. 28, September 20, 1995, 19 pages.

Interprocessor communication overhead is a crucial measure of the power of parallel computing systems—its impact can severely limit the performance of parallel programs. This report presents measurements of communication overhead on three contemporary commercial multicomputer systems: the Intel Paragon, the IBM SP2 and the Meiko CS-2. In each case the time to communicate between processors is presented as a function of message length. The time for global synchronization and memory access is discussed. The performance of these machines in emulating hypercubes and executing random pairwise exchanges is also investigated.

It is shown that the interprocessor communication time depends heavily on the specific communication pattern required. These observations contradict the commonly held belief that communication overhead on contemporary machines is independent of the placement of tasks on processors. The information presented in this report permits the evaluation of the efficiency of parallel algorithm implementations against standard baselines.

ICASE COLLOQUIA
October 1, 1995 - March 31, 1996

Name/Affiliation/Title	Date
Tishkoff, Julian, Air Force Office of Scientific Research "Hypersonic Research Program at AFOSR"	October 6
Sheldon, Frederick, University of Texas at Arlington "Specification and Analysis of Stochastic Properties for Concurrent Systems Expressed using CSP"	October 6
Ristorcelli, Ray, ICASE "Diagnostic Statistics and Heuristic Ideas in the Assessment and Characterization of Complex Turbulent Flows"	October 13
Max, Nelson, Lawrence Livermore National Laboratory "Applications of Texture Mapping to Volume and Flow Visualization"	November 6
Ruge, John, University of Colorado "Algebraic Multigrid (AMG) for Problems in Computational Fluid Dynamics"	November 13
Siva Ram Murthy, C., Indian Institute of Technology (visiting University of Washington) "On Allocation Problems in Parallel and Distributed Computing Systems"	November 27
Chandra, N., FAMU/FSU College of Engineering "Micromechanical Modeling of Finite Deformation Problems in Advanced Materials and Structures"	November 28
Reif, John, Duke University "Predictive Computing: An Emerging Paradigm for Efficient Computation in Dynamic Environments"	November 29
Mukherjee, Amar, University of Central Florida "An Optimal Parallel Algorithm for Volume Ray Casting"	December 4
Eiseman, Peter, Program Development Corporation, New York "GridPro/az3000"	December 6
Keeling, Stephen, Sverdrup Technology, Inc., Arnold Engineering Development Center "Modeling the Chimera Domain Decomposition Approach to Solving Conservation Laws"	December 8

Name/Affiliation/Title	Date
Baden, Scott, University of California, San Diego "Software Infrastructure for Irregular Scientific Computations in Parallel Processors"	December 11
Leutenegger, Scott, University of Denver "Efficient Bulk Loading of R-Trees: A Database Approach for Subset Retrieval of Unstructured CFD Grids"	December 12
Abarbanel, Saul, Tel-Aviv University (on sabbatical at Brown University) "Multi-Dimensional Asymptotically Stable 4th Order Accurate Schemes for the Diffusion Equation"	December 13
Konstantinidou, S., Johns Hopkins University "Communication Issues in Interactive and Dynamic Parallel Computing"	January 18
Johnson, Chris, University of Utah "The SCIRun Computational Steering System"	January 24
Fox, Geoffrey, Syracuse University "HPCC and Web Technologies"	February 1
Jones, Donald, General Motors Research and Development "Engineering Optimization with Response Surfaces"	February 5
Smith, Leslie, Yale University "Two-Dimensional Turbulence and Beyond"	February 6
Brandt, Achi, Weizmann Institute of Science, Israel "Textbook Multigrid Efficiency in CFD"	February 9
Natarajan, Ramesh, IBM TJ Watson Center "An Iterative Scheme for Acoustics Boundary Element Computations"	February 9
Hughes, John, Brown University "User Interfaces for Steady Flow Visualization"	February 20
Mark, Hans, The University of Texas at Austin "The Atomic Bomb: 50 Years Later"	March 4
Mark, Hans, The University of Texas at Austin "The Apollo 13 Accident: A Personal History"	March 5
Whitman, Scott, Cray Research, Inc. "Volume Rendering of Scientific Data on the T3D"	March 7

Name/Affiliation/Title	Date
Painter, James, Los Alamos National Laboratory "Experiences with MPP Rendering on the Cray T3D"	March 8
McClamroch, N. Harris, The University of Michigan "Hybrid Control for Stabilization of a Class of Nonlinear Systems"	March 22
Jameson, Leland, Mitsubishi Heavy Industries, Ltd. "A Wavelet-Optimized, Very High Order Adaptive Grid and Order Numerical Method"	March 26
Masunaga, Akihisa, Mitsubishi Heavy Industries, Ltd. "Application of CFD to the Development of Aerospace Products"	March 28

ICASE STAFF

I. ADMINISTRATIVE

Manuel D. Salas, Director, beginning February 14, 1996. M.S., Aeronautics and Astronautics, Polytechnic Institute of Brooklyn, 1970. Fluid Mechanics and Numerical Analysis.

M. Yousuff Hussaini, Director, through January 3, 1996. Interim Director, through February 13, 1996. Ph.D., Mechanical Engineering, University of California, 1970.

Linda T. Johnson, Office and Financial Administrator

Etta M. Blair, Accounting Supervisor

Barbara A. Cardasis, Administrative Secretary

Shannon L. Keeter, Technical Publications Secretary

Rachel A. Lomas, Payroll and Accounting Clerk

Shelly M. Johnson, Executive Secretary/Visitor Coordinator

Emily N. Todd, Conference Manager

Gwendolyn W. Wesson, Contract Accounting Clerk

Leon M. Clancy, System Manager

Bryan K. Hess, Assistant System Manager

Avik Banerjee, System Operator

II. SCIENCE COUNCIL

Ivo Babuska, Robert Trull Chair in Engineering, The University of Texas-Austin.

Geoffrey Fox, Director, Northeast Parallel Architectural Center, Syracuse University.

Dennis Gannon, Professor, Center for Innovative Computer Applications, Indiana University.

Ashwani Kapila, Professor, Department of Mathematics and Science, Rensselaer Polytechnic Institute.

James P. Kendall, Jet Propulsion Laboratory.

Heinz-Otto Kreiss, Professor, Department of Mathematics, University of California at Los Angeles.

Sanjoy Mitter, Professor of Electrical Engineering, Massachusetts Institute of Technology.

Steven A. Orszag, Professor, Program in Applied & Computational Mathematics, Princeton University.

Paul Rubbert, Unit Chief, Boeing Commercial Airplane Group.

Ahmed Sameh, Department Head of Computer Science, University of Minnesota.

Manuel D. Salas, Director, Institute for Computer Applications in Science and Engineering, NASA Langley Research Center.

III. RESEARCH FELLOWS

Gordon Erlebacher - Ph.D., Plasma Physics, Columbia University, 1983. Fluid Mechanics [Transition and Turbulence]. (November 1989 to November 1999)

Piyush Mehrotra - Ph.D., Computer Science, University of Virginia, 1982. Computer Science [Programming Languages for Multiprocessor Systems]. (January 1991 to September 1999)

IV. SENIOR STAFF SCIENTISTS

Sharath S. Girimaji - Ph.D., Mechanical and Aerospace Engineering, Cornell University, 1990. Fluid Mechanics [Turbulence and Combustion]. (July 1993 to August 1997)

Thomas Jackson - Ph.D., Mathematics, Rensselaer Polytechnic Institute, 1985. Fluid Mechanics. (January 1994 to January 1997)

R. Michael Lewis - Ph.D., Mathematical Sciences, Rice University, 1989. Applied & Numerical Mathematics [Multidisciplinary Design Optimization]. (May 1995 to April 1998)

Josip Loncaric - Ph.D., Applied Mathematics, Harvard University, 1985. Applied & Numerical Mathematics [Multidisciplinary Design Optimization]. (March 1996 to February 1997)

Dimitri Mavriplis - Ph.D., Mechanical and Aerospace Engineering, Princeton University, 1988. Applied & Numerical Mathematics [Grid Techniques for Computational Fluid Dynamics]. (February 1987 to August 1997)

Robert Rubinstein - Ph.D., Mathematics, Massachusetts Institute of Technology, 1972. Fluid Mechanics [Turbulence Modeling]. (January 1995 to January 1998)

David Sidilkover - Ph.D., Applied Mathematics, The Weizmann Institute of Science, 1989. Applied & Numerical Mathematics [Numerical Analysis and Algorithms]. (November 1994 to October 1996)

V. Venkatakrishnan - Ph.D., Mechanical and Aerospace Engineering, Princeton University, 1987. Applied & Numerical Mathematics [Computational Aerodynamics]. (June 1993 to September 1997)

Ye Zhou - Ph.D., Physics, College of William and Mary, 1987. Fluid Mechanics [Turbulence Modeling]. (October 1992 to September 1996)

V. SCIENTIFIC STAFF

Brian G. Allan - Ph.D., Mechanical Engineering, University of California at Berkeley, 1996. Applied & Numerical Mathematics [Multidisciplinary Design Optimization]. (February 1996 to January 1998)

Eyal Arian - Ph.D., Applied Mathematics, The Weizmann Institute of Science, Israel, 1995. Applied & Numerical Mathematics [Multidisciplinary Design Optimization]. (October 1994 to September 1996)

Phillip M. Dickens - Ph.D., Computer Science, University of Virginia, 1992. Computer Science [System Software]. (January 1993 to December 1996)

Stephen Guattery - Ph.D., Computer Science, Carnegie Mellon University, 1995. Computer Science [Parallel Numerical Algorithms, including Partitioning and Mapping]. (September 1995 to September 1997)

M. Ehtesham Hayder - Ph.D., Mechanical and Aerospace Engineering, Princeton University, 1988. Fluid Mechanics [Computational Aeroacoustics]. (September 1995 to September 1997)

Victoria L. Interrante - Ph.D., Computer Science, University of North Carolina at Chapel Hill, 1996. Computer Science [Scientific Visualization]. (March 1996 to February 1998)

Leland M. Jameson - Ph.D., Applied Mathematics, Brown University, 1993. Applied & Numerical Mathematics [Multiresolution Schemes]. (October 1993 to October 1995)

Jim E. Jones - Ph.D., Applied Mathematics, University of Colorado-Boulder, 1995. Computer Science [Parallel Multigrid Methods]. (March 1995 to March 1997)

Kwan-Liu Ma - Ph.D., Computer Science, University of Utah, 1993. Computer Science [Visualization]. (May 1993 to August 1996)

Can Ozturan - Ph.D., Computer Science, Rensselaer Polytechnic Institute, 1995. Computer Science [System Software/Parallel Numerical Algorithms]. (August 1995 to August 1997)

J. Ray Ristorcelli - Ph.D., Mechanical and Aerospace Engineering, Cornell University, 1991. Fluid Mechanics [Turbulence Modeling]. (December 1992 to September 1996)

Moulay D. Tidriri - Ph.D., Fluid Mechanics, University of Paris IX, France, 1992. Computer Science [Domain Decomposition Methods and Parallel Solution of Partial Differential Equations]. (September 1994 to August 1996)

Linda F. Wilson - Ph.D., Electrical and Computer Engineering, University of Texas at Austin, 1994. Computer Science [Investigating the Issues of Object Migration in Object-Oriented Performance Tools and Object-Oriented Modeling of Architectures for Joint Performance/Reliability Analysis]. (August 1994 - July 1996)

VI. SENIOR SYSTEMS ANALYST

Thomas W. Crockett - B.S., Mathematics, College of William & Mary, 1977. Computer Science [System Software for Parallel Computing, Computer Graphics, and Scientific Visualization]. (February 1987 to August 1997)

VII. VISITING SCIENTISTS

Mikhail M. Gilinsky - Ph.D., Aerodynamics, Moscow State University, 1965. Fluid Mechanics [Acoustics and Multidisciplinary Design Optimization]. (March 1996 to July 1996)

Philip Hall - Ph.D., Mathematics, Imperial College, England, 1973. Professor, Department of Mathematics, University of Manchester, England. Fluid Mechanics [Laminar Flow Control]. (August 1995 to January 1996)

Arun K. Somani - Ph.D., Computer Engineering, McGill University, 1985. Professor, Department of Electrical Engineering, University of Washington. Computer Science [Performance Modeling, Reliability Modeling, and Parallel Computing]. (August 1995 to December 1995)

Stephen Woodruff - Ph.D., Aerospace Engineering, The University of Michigan, 1986. Assistant Professor, Division of Engineering, Brown University. Fluid Mechanics. (July 1995 to March 1996)

VIII. SHORT TERM VISITING SCIENTISTS

Maurice Holt - Ph.D., Mathematics, University of Manchester, England, 1948. Professor, College of Engineering, Mechanical Engineering, University of California-Berkeley. Applied & Numerical Mathematics [Godunov Methods]. (February 1996)

Jacques Liandrat - Ph.D., Fluid Mechanics-Numerical Simulations, University of Toulouse, France, 1986. Research Scientist, I.M.S.T., Marseille, France. Applied & Numerical Mathematics. (September 1995 to October 1995)

Bassam Younis - Ph.D., Mechanical Engineering-Fluid Mechanics, Imperial College, London, England, 1984. Senior Lecturer, Department of Civil Engineering, City University, London, England. Fluid Mechanics [Turbulence Modeling in Particular]. (December 1995)

IX. SENIOR RESEARCH ASSOCIATE

David E. Keyes - Ph.D., Applied Mathematics, Harvard University, 1984. Computer Science. [Parallel Numerical Algorithms]

X. CONSULTANTS

Saul Abarbanel - Ph.D., Theoretical Aerodynamics, Massachusetts Institute of Technology, 1959. Professor, Department of Applied Mathematics, Tel-Aviv University, Israel. Applied & Numerical Mathematics [Numerical Analysis and Algorithms]

Ivo Babuska - Ph.D., Technical Science, Technical University, Prague, Czechoslovakia, 1951; Mathematics, Academy of Science, Prague, 1956; D.Sc., Mathematics, Academy of Science, Prague, 1960. Robert Trull Chair in Engineering, TICAM, The University of Texas at Austin. Applied & Numerical Mathematics [Finite Element Methods Associated With Structural Engineering]

Ponnampalam Balakumar - Ph.D., Aeronautics and Astronautics, Massachusetts Institute of Technology, 1986. Associate Professor, Department of Aerospace Engineering, Old Dominion University. Fluid Mechanics [Stability and Transition]

David Banks - Ph.D., Computer Science, University of North Carolina, 1993. Assistant Professor, Department of Computer Science, Mississippi State University. Computer Science [Graphics and Visualization]

H. Thomas Banks - Ph.D., Applied Mathematics, Purdue University, 1967. Professor, Department of Mathematics, Center for Research in Scientific Computations, North Carolina State University. Applied & Numerical Mathematics [Control Theory]

Richard W. Barnwell - Ph.D., Engineering Mechanics, Virginia Polytechnic Institute and State University, 1968. Professor, Department of Aerospace and Ocean Engineering, Engineering Science and Mechanics. Virginia Polytechnic Institute and State University. Fluid Mechanics [Turbulence Modeling]

Alvin Bayliss - Ph.D., Mathematics, New York University, 1975. Associate Professor, Technological Institute, Northwestern University. Fluid Mechanics [Numerical Solution of the Equations of Fluid Flow and Acoustics]

John A. Burns - Ph.D., Mathematics, University of Oklahoma, 1973. Professor, Virginia Polytechnic Institute and State University. Applied & Numerical Mathematics [Numerical Methods in Optimal Design and Control]

Xiao-Chuan Cai - Ph.D., Mathematics, Courant Institute of Mathematical Science, New York University, 1989. Assistant Professor, Department of Computer Science, University of Colorado-Boulder. Computer Science [Numerical Analysis and Parallel Computing]

Tony F. Chan - Ph.D., Computer Science, Stanford University, 1978. Professor, Department of Mathematics, University of California-Los Angeles. Applied & Numerical Mathematics [Multigrid and Domain Decomposition for Unstructured Grid]

Barbara M. Chapman - M.S., Mathematics, University of Canterbury, Christchurch, New Zealand, 1985. Director, European Institute for Parallel Computing, University of Vienna. Computer Science [Parallel Language Extensions and Optimizations for Parallel Compilers]

Gianfranco Ciardo - Ph.D., Computer Science, Duke University, 1989. Assistant Professor, The College of William & Mary. Computer Science [Reliability Models]

Ayodeji O. Demuren - Ph.D., Mechanical Engineering, Imperial College London, United Kingdom, 1979. Associate Professor, Department of Mechanical Engineering and Mechanics, Old Dominion University. Fluid Mechanics [Numerical Modeling of Turbulent Flows]

Geoffrey Fox - Ph.D., Physics, Cambridge University, 1967. Professor, Department of Computer Science, Syracuse University. Computer Science [Networking]

Dennis B. Gannon - Ph.D., Mathematics, University of California, Davis, 1974. Professor, Department of Computer Science, Indiana University-Bloomington. Computer Science [Investigation of Algorithms and Programming Techniques for Parallel Computers]

James F. Geer - Ph.D., Applied Mathematics, New York University, 1967. Professor, Systems Science and Mathematical Sciences, Watson School of Engineering, Applied Science and Technology, SUNY-Binghamton. Applied & Numerical Mathematics [Approximation Method of Solutions for Partial Differential Equations]

David Gottlieb - Ph.D., Numerical Analysis, Tel-Aviv University, Israel, 1972. Professor, Division of Applied Mathematics, Brown University. Applied & Numerical Mathematics [Boundary Conditions for Hyperbolic Systems]

Chester E. Grosch - Ph.D., Physics and Fluid Dynamics, Stevens Institute of Technology, 1967. Professor, Department of Computer Science and Slover Professor, Department of Oceanography, Old Dominion University. Fluid Mechanics [Acoustics]

Max Gunzburger - Ph.D., Applied Mathematics, New York University, 1969. Professor, Department of Mathematics, Iowa State University. Applied & Numerical Mathematics [Numerical Methods for Flow Control Problems]

Matthew D. Haines - Ph.D., Computer Science, Carnegie Mellon University, 1995. Assistant Professor, Department of Computer Science, University of Wyoming. Computer Science [Parallel Programming Environment and Run Time Systems]

Gene J.-W. Hou - Ph.D., Computational Mechanics, Design Optimization, University of Iowa, 1983. Associate Professor, Mechanical Engineering Department, Old Dominion University. Applied & Numerical Mathematics [Computational Mechanics Design Optimization]

Fang Q. Hu - Ph.D., Applied Mathematics, Florida State University, 1990. Assistant Professor, Department of Mathematics and Statistics, Old Dominion University. Fluid Mechanics [Instability and Transition]

Hideaki Kaneko - Ph.D., Applied Functional Analysis, Clemson University, 1979. Professor, Department of Mathematics and Statistics, Old Dominion University. Applied & Numerical Mathematics [Multidisciplinary Design Optimization]

Ashwani K. Kapila - Ph.D., Theoretical and Applied Mechanics, Cornell University, 1975. Associate Professor, Department of Mathematical Sciences, Rensselaer Polytechnic Institute. Fluid Mechanics [Mathematical Combustion]

Ken Kennedy - Ph.D., Computer Science, New York University, 1971. Chairman, Department of Computer Science, Rice University. Computer Science [Parallel Compilers and Languages]

Charles Koelbel - Ph.D., Computer Science, Purdue University, 1990. Research Scientist, Department of Computer Science, Rice University. Computer Science [Compilers for Parallel Computers]

David A. Kopriva - Ph.D., Applied Mathematics, University of Arizona, 1982. Associate Professor, Department of Mathematics and SCRI, Florida State University. Applied & Numerical Mathematics [Development and Implementation of Spectral Methods for Viscous Compressible Flows]

Frank Kozusko - Ph.D., Computational and Applied Mathematics, Old Dominion University, 1995. Assistant Professor, Department of Mathematics, Hampton University. Fluid Mechanics

Heinz-Otto Kreiss - Ph.D., Mathematics, Royal Institute of Technology, Sweden, 1960. Professor, Department of Applied Mathematics, California Institute of Technology, Applied & Numerical Mathematics [Numerical Solution of Partial Differential Equations]

David G. Lasseigne - Ph.D., Applied Mathematics, Northwestern University, 1985. Assistant Professor, Department of Mathematics and Statistics, Old Dominion University. Fluid Mechanics [Asymptotic and Numerical Methods for Computational Fluid Dynamics]

Scott T. Leutenegger - Ph.D., Computer Science, University of Wisconsin-Madison, 1990. Assistant Professor, Department of Mathematics and Computer Science, University of Denver. Computer Science [System Software Related to Databases for Scientific Data]

Robert W. MacCormack - M.S., Mathematics, Stanford University. Professor, Department of Aeronautics and Astronautics, Stanford University. Fluid Mechanics [Numerical Methods for the Solution of the Equations of Fluid Mechanics]

Kurt Maly - Ph.D., Computer Science, Courant Institute, New York University, 1973. Kaufman Professor and Chair, Department of Computer Science, Old Dominion University. Computer Science [High Performance Communication]

James E. Martin - Ph.D., Applied Mathematics, Brown University, 1991. Assistant Professor, Department of Mathematics, Christopher Newport University. Fluid Mechanics [Turbulence and Computation]

Sanjoy K. Mitter - Ph.D., Electrical Engineering, Imperial College of Science & Technology, London, 1965. Professor of Electrical Engineering, Co-Director, Laboratory for Information and Decision Systems, Director, Center for Intelligent Control Systems, Massachusetts Institute of Technology. Fluid Mechanics [Control Theory]

David M. Nicol - Ph.D., Computer Science, University of Virginia, 1985. Professor, Department of Computer Science, College of William & Mary. Computer Science [Mapping Algorithms onto Parallel Computing Systems]

R.A. Nicolaides - Ph.D., Computer Science, University of London, 1972. Professor, Department of Mathematics, Carnegie Mellon University. Applied & Numerical Mathematics [Numerical Solution of Partial Differential Equations]

Demetrius Papageorgiou - Ph.D., Mathematics, University of London, 1985. Assistant Professor, Department of Mathematics, New Jersey Institute of Technology. Fluid Mechanics [Theoretical and Computational Fluid Dynamics]

Alex Pothén - Ph.D., Applied Mathematics, Cornell University, 1984. Professor, Department of Computer Science, Old Dominion University. Computer Science [Parallel Numerical Algorithms]

Daniel A. Reed - Ph.D., Computer Science, Purdue University, 1983. Assistant Professor, Department of Computer Science, University of Illinois. Computer Science [Performance Evaluation of Computer Systems]

John H. Reif - Ph.D., Applied Mathematics, Harvard University, 1977. Professor, Department of Computer Science, Duke University. Computer Science [Parallel Algorithms]

Philip L. Roe - Ph.D., Aeronautics, University of Cambridge, United Kingdom, 1962. Professor, Department of Aerospace Engineering, University of Michigan. Applied & Numerical Mathematics [Numerical Mathematics and Aeroacoustics]

Ahmed H. Sameh - Ph.D., Civil Engineering, University of Illinois, 1968. Head, William Norris Chair, and Professor, Department of Computer Science, University of Minnesota. Computer Science [Numerical Algorithms]

Ralph C. Smith - Ph.D., Mathematics, Montana State University, 1990. Assistant Professor, Department of Mathematics, Iowa State University. Applied & Numerical Mathematics [Optimal Control Techniques for Structural Acoustics Problems]

Charles G. Speziale - Ph.D., Aerospace and Mechanical Sciences, Princeton University, 1978. Professor, Aerospace & Mechanical Engineering Department, Boston University. Fluid Mechanics [Turbulence Modeling]

Shlomo Ta'asan - Ph.D., Applied Mathematics, The Weizmann Institute of Science, 1985. Professor, Department of Mathematics, Carnegie Mellon University. Applied & Numerical Mathematics [Numerical Analysis and Algorithm Development]

Siva Thangam - Ph.D., Mechanical Engineering, Rutgers University, 1980. Professor, Department of Mechanical Engineering, Stevens Institute of Technology. Fluid Mechanics [Turbulence Modeling and Simulation]

Lu Ting - Ph.D., Aeronautics, New York University, 1951. Professor, Courant Institute of Mathematical Sciences, New York University. Fluid Mechanics [Nonlinear Acoustic/Structure Interaction Problems]

Virginia Torczon - Ph.D., Mathematical Sciences, Rice University, 1989. Assistant Professor, Department of Computer Science, The College of William & Mary. Computer Science [Parallel Algorithms for Optimization including Multidisciplinary Design Optimization]

Kishor Trivedi - Ph.D., Computer Science, University of Illinois-Urbana, 1974. Professor, Department of Electrical Engineering, Duke University. Computer Science [Performance and Reliability Modeling Methods, Tools and Applications]

George M. Vahala - Ph.D., Physics, University of Iowa, 1972. Professor, Department of Physics, The College of William & Mary. Fluid Mechanics [Group Renormalization Methods for Turbulence Approximation]

Bram van Leer - Ph.D., Theoretical Astrophysics, Leiden State University, The Netherlands, 1970. Professor, Department of Aerospace Engineering, University of Michigan. Applied & Numerical Mathematics [Computational Fluid Dynamics]

Hans Zima - Ph.D., Mathematics, University of Vienna, Austria, 1964. Professor, Institute for Software Technology and Parallel Systems, University of Vienna, Austria. Computer Science [Compiler Development for Parallel and Distributed Multiprocessors]

XI. GRADUATE STUDENTS

Abdelkader Baggag - Graduate Student at The University of Minnesota. (September 1995 to Present)

David C. Cronk - Graduate Student at The College of William & Mary. (August 1993 to Present)

Dawn M. Galayda - Graduate Student at The College of William & Mary. (September 1995 to Present)

Nilan Karunaratne - Graduate Student at Auburn University. (August 1995 to Present)

Joe L. Manthey - Graduate Student at Old Dominion University. (September 1993 to Present)

Debora Pilkey - Graduate Student at Virginia Polytechnic Institute and State University. (October 1995 to Present)

Kevin Roe - Graduate Student at Syracuse University. (May 1995 to Present)

Robert V. Wilson - Graduate Student at Old Dominion University. (October 1992 to Present)

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE May 1996	3. REPORT TYPE AND DATES COVERED Contractor Report		
4. TITLE AND SUBTITLE Semiannual Report. October 1, 1995 through March 31, 1996		5. FUNDING NUMBERS C NAS1-19480 WU 505-90-52-01		
6. AUTHOR(S)				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Institute for Computer Applications in Science and Engineering Mail Stop 132C, NASA Langley Research Center Hampton, VA 23681-0001		8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) National Aeronautics and Space Administration Langley Research Center Hampton, VA 23681-0001		10. SPONSORING/MONITORING AGENCY REPORT NUMBER NASA CR-198338		
11. SUPPLEMENTARY NOTES Langley Technical Monitor: Dennis M. Bushnell Final Report				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Unclassified-Unlimited Subject Category 59		12b. DISTRIBUTION CODE		
13. ABSTRACT (Maximum 200 words) This report summarizes research conducted at the Institute for Computer Applications in Science and Engineering in applied mathematics, fluid mechanics, and computer science during the period October 1, 1995 through March 31, 1996.				
14. SUBJECT TERMS Applied Mathematics; Numerical Analysis; Fluid Mechanics; Computer Science			15. NUMBER OF PAGES 68	
			16. PRICE CODE A04	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT	